Pund-IT

PUND-IT, INC.

Weekly Review

MAY 3, 2006

Getting Back to Basics on Data Security

By David G. Hill, The Mesabi Group

Security, in general, and data security, in particular, is receiving increased attention from both IT organizations and vendors. On the data security side, the much publicized string of disappearances and/or thefts of tape media containing personal sensitive information on individuals, such as credit card and social security numbers, have raised a hue and cry about data confidentiality. But data confidentiality is only part of data security.

In the way that we typically think about the subject, data security is part of data protection, but is not synonymous with data protection. Data security focuses on defenses to threats that are willful; that is threats of human origin, whether or not they are intentional or unintentional. For example, data security is concerned with stopping viruses that could corrupt data, but not with the corruption of a database table by an application. Yes, the result would be the same — corrupted data — but a security professional would deal with the virus, while a database professional would deal with the database corruption. Similarly, from a physical perspective, the lack of a lock on a data center door would be a security issue; the fire suppressant system in the data center is not, but *is* a data protection issue.

All four objectives of data protection — data preservation, data availability, data responsiveness, and data confidentiality — apply to data security, but the two principal objectives are data preservation and data confidentiality.

Data Preservation Is the First Goal of Data Security

When seeking data security, a user's first focus has to be on data preservation. Data preservation can be divided into two parts — data integrity and data survival. Data integrity means that the data retrieved from storage is the same data that was put in; the bit patterns are the same and all the bits are there (i.e., data completeness). Data integrity also means that the data has not been subject to unauthorized modification since its creation or since the last authorized change. If an unauthorized modification has taken place, the result is either corruption — the data is unusable — or potentially inaccurate, misleading, or false; an alteration of the data that deceives the user or a user application.

Data survival means that the data is there in the first place and can be found. That is, the data has not been subject to unauthorized destruction or that linkages which enable the data to be retrieved have not been disabled or deleted. "Death" of data may lead to simple malfunctioning business processes (such as the inability to invoice a customer) or the significant inability to furnish evidence (which could lead to serious economic consequences in a legal case). Data preservation is therefore an essential component of data security (and of data protection). Preserving your data helps you preserve your business.

Data preservation is essential for meeting compliance requirements. Data preservation practices for non-compliant data are at the discretion of the organization that owns the data. That is not true for compliant data — compliant data is subject to data retention requirements that may be set by regulators and/or government agencies. The organization that owns the compliant data has a choice in which policies, practices, procedures, and technologies to use to enforce data preservation requirements; it does not have the choice to do nothing.

Confidentiality Is a Public Concern about Private Information

Confidentiality is also an essential goal of data security. Confidentiality is the prevention of unauthorized disclosure of information. However, whereas data *preservation* focuses on the *intrinsic* value of information to those authorized to use it; *confidentiality* focuses on the *extrinsic* value of that information to unauthorized parties.

The key issue with confidentiality has not been the unauthorized disclosure of proprietary information, such as trade secrets or non-consumer customer lists. If the owners of proprietary information fail to protect the confidentiality of that information adequately, then that organization alone (along with its stockholders) suffers the consequences. That is a private not a public concern from a governmental perspective. Rather, the issue with confidentiality has been the unauthorized disclosure of private information about individuals — such as social security numbers and credit card numbers — either intentionally or unintentionally by organizations that have possession of that information for legitimate purposes— such authorizers of credit card transactions.

Unless there are externally imposed consequences on those organizations (such as result of a lawsuit), possessors of other individuals' private information do not suffer from the exposure of that information to unauthorized third parties. However, individuals whose information has been exposed can suffer losses ranging from a loss of privacy (medical records) to economic loss (identity theft).

Governments continue to create legislation that attempts to correct that risk imbalance between the possessor of an individual's private information and the individual. Possessors of private information now have to contend with laws and with threats of litigation. One of the consequences is public exposure of the failure to protect the confidentiality of private individual information. That can result not only in public embarrassment, but also in a possible negative impact on the brand or market value of a firm. Consequently, organizations are giving much more consideration to confidentiality policies and practices today than they have in the past.

The Special Case of Storage Security

We often talk about security from the infrastructure perspective — that is network security or storage security. However, when the talk focuses on storage security, the concern generally is not of physical security (although there has to be some such as a lock on the door to the data center). After all, when a tape cartridge containing confidential data is stolen or lost, no one worries about the replacement cost of the lost tape. No, storage security is typically about *logical* storage security, which is simply concern about the data itself, so is the discussion really does focus on data security. All that is different is the perspective. Yet that perspective is critically important.

In the days when all disk storage was directly attached to the server and all the I/Os flowed between the two, storage security was a non-issue. When Fibre Channel storage

area networks came into existence, the presumption was that the SAN was secure. After all, it was a physically-secure network within a data center that required specialized knowledge to access. Now administrators access SAN storage not only within data centers, but also through a switched long-latency network called the Internet. Storage administrators can make changes to their SAN from a Web browser in the convenience of their own home if necessary. As a result, storage technically and practically is no longer secure.

Moreover, the rising awareness is that FC SANs are not as secure as they originally were thought to be. SANs have had only "security through obscurity." That means that SANs have generally been safe only because the security holes have not been highly publicized or accessible. While those potential holes are not likely to make the front pages, they are well-documented in publicly-accessible materials.
FC SANs were not designed with security features that are common for IP networks. For example, FC lacks certain access control features that security professionals consider necessary for proper levels of authentication. NAS and iSCSI cannot criticize FC SANs for lack of security, as they have their own structural security defects.

Putting Up a Good Defense

An IT organization should look at data security problems from a global infrastructure perspective — server security (e.g., operating system and application), network security, and storage security. These security defenses focus on the logical value of the data — as before, if a magnetic tape is lost or stolen, no one is concerned about the cost of the tape, only the value of the information. And the defenses focus on the data: how servers process data, how networks move data, and how storage stores data.

All of these defenses tie together in layers, i.e., fallback positions. Security discussions often revolve about such defense-related terms as threat vectors, attack surfaces, and perimeter defenses. How well data should be protected depends upon its value — and different data has different value — as well as the cost of protecting that data. Security costs money; what can an organization afford — the safe that is rated for 15 minutes of protection, one that is rated for 30 minutes, or a bank vault?

Another key principle should be limiting the amount of damage any one individual — either an outsider or an insider — can cause either deliberately or inadvertently. That is especially critical when Web browsers are involved in management processes, as insiders can be outside the firewalls. Even if authentication credentials are submitted by the authorized person (i.e., not stolen), the privileges allowed in such circumstances should be limited. For example, no remote superuser privileges should be permitted. That is not a question of trust in or the competence of a particular individual, but rather a sensible general policy to reduce exposure. Of course, among those actions necessary are that event logs should be compiled for auditing purposes and that clear policies be designed and enforced to limit the amount of damage.

Storage Vendors Are Getting Serious about Data Security

That data security remains a serious problem is the bad news. The good news is that the situation has caught the attention of some storage vendors, including those who have the incentives and resources to address the situation. That is not to say that a number of technologies and approaches to data security are not already available today. They are, and organizations can take advantage of what is already available to improve their data security. However, there is more to be done.

EMC recently announced a comprehensive strategy that revolves around information security as an information management problem throughout the life-cycle of information. An EMC Assessment Service for Storage Security based on the National Security Agency's Information Assurance Methodology is a professional service offering that is only the starting point. EMC is following a product security scorecard approach to make sure that its products meet storage security compliance requirements. Overall, EMC plans to invest between 5% and 10% of its ongoing R&D budget dollars on security-related activities. Since EMC's R&D budget is $1.2B that represents a significant number. EMC means business.

NetApp has its own Uncompromised Security Initiative. NetApp feels that the current focus on network perimeters and on specific devices and people fails to consider data throughout its lifecycle. Moreover, enterprises cannot trade off performance, interoperability, high availability, and simplicity when they deploy security solutions — hence the name uncompromised security. NetApp also means business.

Sun focuses a great deal of company time and resources on identity management. Sun views the future as identity-managing everything. Identity management is more than security; it is also the future of what Sun calls integrated data management. Identity management associates users with data so that security and service levels can be tracked while business processes are taking place. Sun means business, too.

Those are the announced programs. Rest assured that HP, IBM, and other major players probably have something up their sleeves, as well. And naturally there are smaller companies who focus on data security — Decru (now part of NetApp), NeoScale, and SenSage among them.

Mission Accomplished?

Ah, no. Yes, vendors are doing all this work for the best interests of their customers. But they are also doing it because it is in their best interests, too. In the near future some vendors may be able to differentiate themselves based on their storage security solutions. And that is likely to result in both significant hype and substance, as along with the need to carefully discern which is which.

Since the vendors are coming at the subject from a number of different angles, process of separating the wheat (substance) from the chaff (hype) may be easier said than done. Therefore, going back to the basics — data preservation and data confidentiality — and doing what is necessary to ensure each will be critical in the process of evaluating, choosing, and deploying a successful data security solution.

*David G. Hill is principal of the **Mesabi Group** (www.mesabigroup.com). The Mesabi Group focuses on the revolutions in Storage Networking and Storage Management, and helps clients make the best and most efficient use of information for business value.*