ISSUE PAPER

Information Risk Management in the Enterprise

John Burke, Principal Research Analyst

Executive Summary

Enterprise IT security is being pulled steadily towards a risk-based view of the world. Companies need to understand their tolerance for risk, and embrace technologies and practices that allow them to meet, but not exceed, that tolerance. The disciplines of information stewardship provide a lens through which the enterprise can focus its actions in information risk management. By focusing on the discipline of information protection, it can choose where and how to apply technologies, such as encryption, to maximize the return on risks of information leak or theft. Focusing on data quality management can minimize both the operational risks from inconsistent or incorrect data, and the legal risks from lapses in compliance, inadvertent disclosure, or unintentional failure to disclose information in court. Focusing on continuity mitigates risk from data being unavailable due to natural disaster, systems break down, or attack.

Growth of a Risk Management Perspective in Enterprise IT Security

In Nemertes *Security and Information Protection* research benchmark, fully two-thirds of participants reported that they have decided for security reasons not to use a technology or service in which users or business lines are interested. Typically, these were productivity-enhancing products or applications, including VOIP, collaborative tools or mobile/wireless networks or services.

This is reasonable from IT's traditional security perspective. IT security teams are set up to prevent and react to security problems, not to set acceptable levels of risk. They evaluate a new application against the goal of the best possible security, and if it introduces significant new levels of risk, try to quash it.

Such an approach is deeply shortsighted from an enterprise perspective. There is, after all, a hidden cost—an opportunity cost—to avoiding useful technologies for security reasons.

Note also that companies pay either way: they either invest up front in a more secure infrastructure that enables them to deploy newer technologies with confidence, or pay the price of non-deployment in the form of reduced revenues and competitiveness. This tradeoff is not obvious from the standpoint of the IT department: Security investments come out the IT budget, but the cost of not investing is borne by others. This makes the avoidance tactically sound for the IT Department but strategically unsound for the enterprise.

More sound is an assessment of the risks of the new technology versus the rewards its adoption could bring, and efforts to mitigate if not completely eliminate that risk should adoption proceed. This theme of "security as enabler via strategic risk mitigation" is part of an emerging trend in IT – a shift from a threat/response perspective to a risk and reward perspective.

At the same time, the enterprise focus on the security of information resources is skyrocketing in the wake of ever-increasing numbers of major data leaks and in the face of a professional, profit-seeking, market-driven world of cybercriminals. Information risk management therefore becomes a new and major focus of enterprise activity, and is best understood in the context of the larger issue of enterprise information stewardship.

Stewardship and Risk

More important than any given technology in the management of enterprise information are the perspectives guiding strategy, policy, process, and systems architecture.

Information stewardship is a key perspective Nemertes highlights. Information stewardship calls for holistic data management in the enterprise: defining and enforcing policy to guide the acquisition, management, and storage lifecycle of data, and the protection of data from theft, leak, or disaster. Broadly speaking, information stewardship includes data quality management, information lifecycle management, business continuity planning, information protection, and compliance. Nemertes research shows that enterprises that manage these intertwined issues as a set are more successful dealing with them than those that treat them as disjoint.

It is straightforward to re-frame the major components of information stewardship as risk management activities. (Please see Figure 1: Information Stewardship Disciplines and Information Risks, Page 3.) After all, the clear underlying reason for something like business continuance planning is to mitigate the risk of an IT service outage hampering business operations; for compliance, the risk of adverse consequences from failing to obey the law or meet industry standards; for information protection, the risk of data being stolen or leaked. Data quality management – reducing the amount of garbage in enterprise systems, to reduce the amount of garbage coming out of them – also

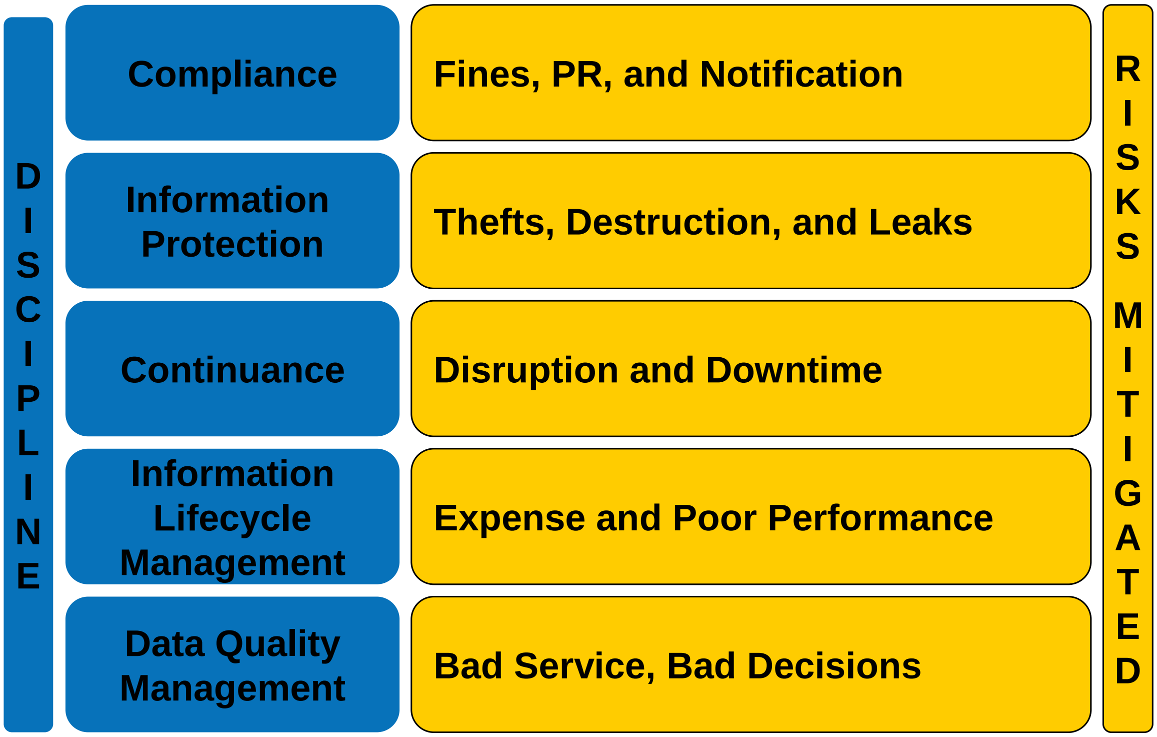

DISCIPLINE	RISKS MITIGATED
Compliance	Fines, PR, and Notification
Information Protection	Thefts, Destruction, and Leaks
Continuance	Disruption and Downtime
Information Lifecycle Management	Expense and Poor Performance
Data Quality Management	Bad Service, Bad Decisions

Figure 1: Information Stewardship Disciplines and Information Risks

reduces risk, including the risk of providing bad service or making bad decisions based on bad data. Information lifecycle management optimizes the balance between two complementary risks: the risk that IT will waste resources (time and money) by using the most expensive kinds of storage for data not requiring it; and conversely that data will be not be available as quickly as needed when it is needed. Business continuity planning guards against the risk of information being entirely inaccessible.

Security and Information Protection

Information protection, reducing the risk to the enterprise of information being stolen or leaked, is a hot concern amongst enterprises, the objective of 38.6% top security projects in 2007 and of 40% planned for 2008 according to Nemertes *Security and Information Protection* benchmark.

Information protection encompasses primarily technologies such as network and storage encryption, and enterprise rights management. Top-most on IT's mind this year is protecting the data residing on enterprise laptops, and encryption is the tool of choice. Solutions range from self-encrypting hard drives to freeware storage encryption to commercial products, and about 10% of participants already have something deployed or in deployment; more than twice that many are evaluating their options, at the laptop, desktop, and server. At issue are not just the specific technologies available but also the degree to which the enterprise is exposed to risk by information exposure or theft, and the rewards associated with that data being freely available on internal systems, desktops, and laptops.

Every business continuously balances risk and reward to find a way to achieve the best returns at an acceptable level of risk. For IT security professionals, this is the most difficult part of the job: objectively analyzing risk *in the context of the business goals and possible return on investment.* It may seem counterintuitive, but the end-goal for the business as a whole is not to achieve zero risk—shutting down would be the best way to achieve that, just as the server least vulnerable to attack is the server that is not turned on. Rather, the business goal is to allow the *maximum acceptable* level of risk – to live at the limit of the organization's "risk tolerance." Every business decision is about risk—getting the maximum return for a given level of risk; IT decisions are no different.

Increasingly, companies are recognizing that the CSO's central and fundamental job revolves not around technologies or even polices but around evaluating the risk of different business choices and then directing the appropriate mitigation strategy. In other words, the CSO is tantamount to a strategic information risk manager, looking not just at technology in isolation. The CSO's ideal metric is therefore not return-on-investment (ROI) or total-cost-of-ownership (TCO) but instead return-on-risk (ROR) or total-risk-of-implementation (TRI).

It's important to note that the CSO's role is to manage risk, not to set acceptable risk levels – that's the job of the CEO and/or board of directors. It is also important to see that this role creates a potential conflict of interest between the CSO and the CIO, whose job is to maximize ROI—to invest in technologies that deliver the maximum bang for the buck. If the CIO can override the CSO's decisions – as when a CSO reports to a CIO – the CSO's ability to deliver ROR is compromised and the company's ability to successfully manage risk is threatened. Consequently, we see the CSO becoming an officer of the company at growing number of companies, as enterprises acknowledge that the CSO is really a peer and counterpart of the CIO.

Compliance and Auditing

From the emerging perspective of risk management, the various threats posed by information loss or leakage that IT secures against all represent risks to the enterprise. These may be loss of competitive advantage following theft of intellectual property; risk of disrupted business; or risk of business lost or never won because of bad PR following a major compromise.

The presence of regulatory requirements for accountability and accuracy in various aspects of enterprise information management adds weight to such risk assessments by increasing the degree and type of costs associated with breaches. Requirements for notification of breaches do so as well: public disclosure increases the chance of losing business (since more people will know of the breach), and direct notification of concerned parties layers on the hard-dollar costs – currently averaging $80 to $88 per record compromised – resulting from the need to find, notify, and compensate victims of identity theft.

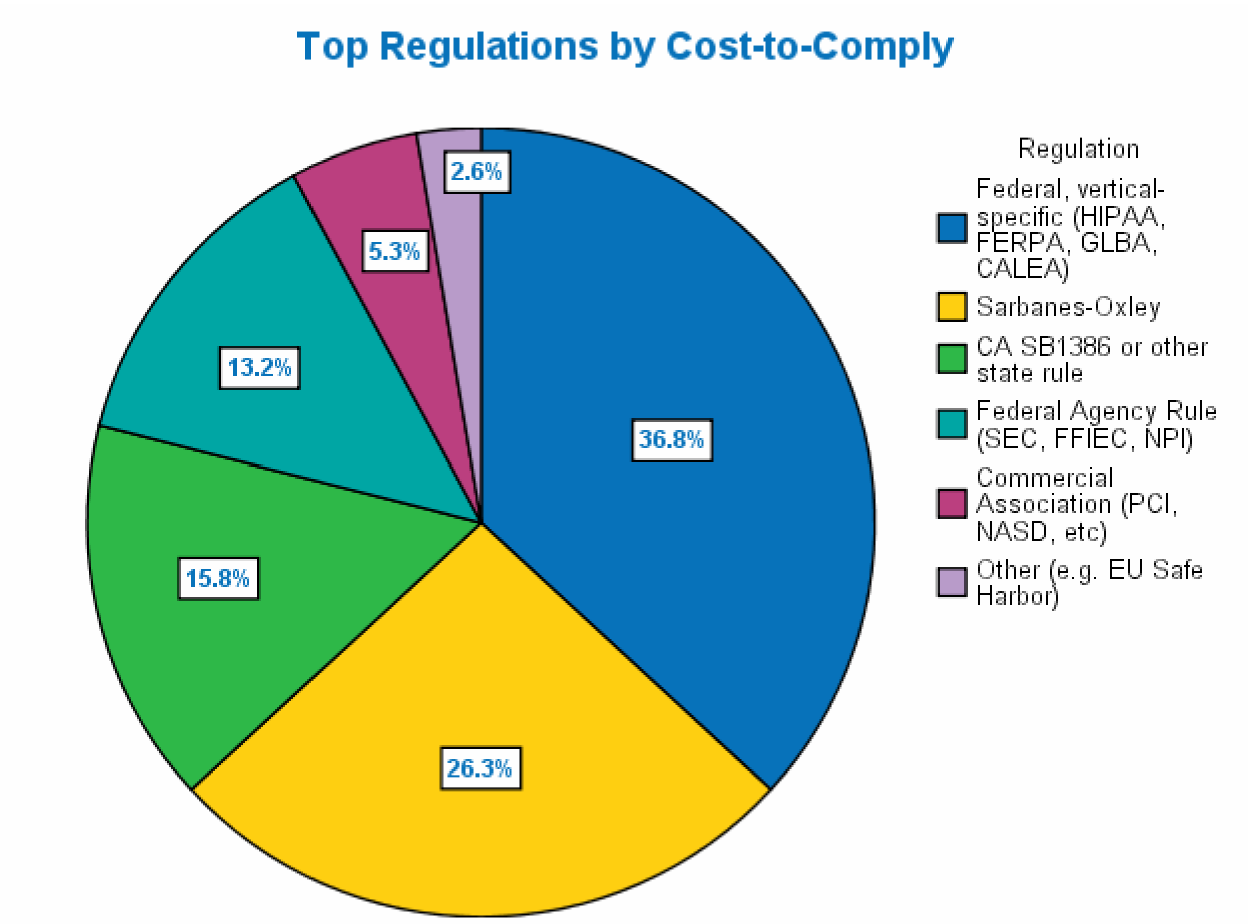

Figure 2: Cost of Compliance, by Regulation

Compliance requirements, as a result, increasingly influence risk analyses and drive security activity, planning, and spending. Also, as the security team slowly morphs into part of the "corporate risk mitigation" team, challenges previously handled by legal and compliance staffs are also starting to fall into its purview. Over half of Nemertes *Security and Information Protection* benchmark participants say they're responsible for e-discovery (55.6%) and compliance (81.5%).

Although the compliance requirements on an enterprise may be many, varied, and complex, IT executives are acutely aware of the parts of the regulatory landscape that are costing them the most. (Please see Figure 2: Cost of Compliance, by Regulation, Page 5). Somewhat surprisingly, despite the volume and intensity of griping about it, the Sarbanes-Oxley Act (SOX) is not the most onerous regulation; rather, that honor goes to the various vertical-specific federal regulations such as the Health Information Portability and Accountability Act (HIPAA) or the Graham-Leach-Blighly Act (GLBA). A solid majority (60%) of participants expect their compliance spending to increase, and the remainder expect it to stay level. No one expects compliance costs to go down as changes in the regulatory landscape (typically new requirements whether in law or professional codes) absorb any savings garnered through automation or improvements in tools.

Closely associated with compliance is auditing against defined controls intended to implement that compliance. The significance of audits is multi-faceted. Of course, some regulatory regimes, like the PCI's, mandate them, and so they represent "table stakes" in some lines of business. They also provide a backstop to IT when it is impossible or impractical to implement practices, such as full separation of duties that preclude some forms of misfeasance and malfeasance. Last, in an IT organization striving to continuously improve its security profile and execution—in security shops with mature metrics for success—audits provide the second-best feedback on how well IT is executing compared to policy and plan. (The best feedback comes from post-mortems on actual compromises.)

Nearly three quarters—71.1%—of participants in *Security and Information Protection* conduct internal audits, and they most often do them annually (54%) or quarterly (25%). Annual audits are usually comprehensive, while more frequent ones tend to cover only a subset of policies, procedures or systems. Slightly fewer participants (just shy of 66%) conduct regular external audits. Of these, the great majority (nearly three quarters) are audited annually, though over 10% have frequent, ad-hoc external audits, whether they want them or not, courtesy of regulating bodies or corporate HQ.

Audits require raw materials, and in addition to discussions with humans and spot checks on systems configurations, auditing for compliance relies heavily on auditing logs. Audit trails in the form of logs of network, system and application activities are a major means of demonstrating compliance. This, along with the increasing sophistication of log analysis for security, has driven the rapid adoption of log aggregation: about 64% of participants collected logs from many sources and aggregated them for analysis, reporting, and retention. Just under half of that group aggregate all infrastructure-related logs—server, router, firewall, etc.—and some even collect desktop logs as well. About 78% of participants archive some or all logs they collect. This can equate to a phenomenal amount of disk space, even if a log-normalizing system is in place to reduce the amount of duplicative information to a bare minimum. This information, however, is rarely used outside the log management and security purposes for which it is initially collected.

Discovery and Integrity

Discovery is the process of producing information for use in a lawsuit, whether in defense of the enterprise or at the request of someone making a claim against it. Discovery is closely related to compliance in IT thinking, and as noted above is increasingly a responsibility of IT. The recently revised rules for e-discovery—discovery of electronically stored information (ESI)—ease some of the burdens of managing discovery while pushing organizations to improve their enterprise information stewardship overall.

The biggest and most important change has been to recognize ESI as a separate category in discovery, distinct from documents and objects. With that distinction understood, many other rules were amended to take into account the speed and volume of response possible with automated production of electronic information, as well as the varying degrees of recoverability of such data and the ephemeral nature of much information in systems.

Organizations creating e-discovery plans or responding to information requests will have an easier time now—as long as they have implemented good information stewardship practices in advance! Requirements to describe the ESI they will use on their own behalf bolster the case for data classification and strong data management. Specific protections from risk of sanction for “normal course of business” deletion of data, and from the risk posed by “inadvertent production” of information that should have been protected from discovery have been created, but require “good faith” in information handling. Demonstrating that you are managing and producing information in good faith is easier if you have been following the basic tenet of information stewardship: for all information in the enterprise, define and enforce policy governing all access to it and its entire lifecycle.

Data classification is fundamental to robust compliance and e-discovery support (and to information lifecycle management, continuity planning, and security). Inadvertent production in e-discovery and other forms of unintentional release of information are often the result of data being misclassified or misunderstood.

A related problem, cited by several participants in Nemertes *Services-Oriented Architectures and Applications* benchmark, arises as IT uses SOA to integrate previously separate silos of functionality. Dissolving silos reveal the fact that different departments or applications disagree on data meaning or data values for fields they must share, an enterprise’s master data. About 25% of participants use tools to help with master data management (MDM), and nearly 53% of the rest say they plan to.

Master data management is one aspect of enterprise data quality management (DQM), the information stewardship discipline aimed at ensuring that the mission-critical data within an enterprise is reliable, accurate and complete. This emphasis on integrity is increasingly important as data is used by more people to make more decisions within the enterprise, and as compliance requires that certain types of data be accurate during the entire course of its life. Moreover, e-discovery tools make it possible to hand over ever more data ever more quickly in legal proceedings, with the resulting risk of handing over more than is intended. That risk is compounded by the fact that enterprises are keeping data around longer than ever, both for operational reasons and for legal ones. In fact, it is increasingly popular to plan to retain some kinds of information “forever” against the possibility that even after legal or operational requirements have passed, there is value in holding it against possible discovery requests. The thinking is that it is better to have information to bring to bear than to be at the mercy of whatever information opponents in a proceeding might have. Keeping data accurate and in agreement everywhere is one way to both prevent inadvertent disclosure and enable retrieval of all useful information in the face of this continuing, mushrooming growth.

Availability

The integrity information is, to some degree, moot if the information is not available for use. In the past, IT folks framed business continuity as being somewhat synonymous with disaster recovery: planning for disastrous disruptions to the IT infrastructure or staff, with time lines from minutes to days for restoration of services. The new understanding of business continuity emphasizes the "continuity" and implies planning to ensure that a company's critical online processes stay available regardless of what's happened to the infrastructure or people. The old goal – for periods of *normal* operation – was uptime on par with the telephone system: 99.999%, the now canonical "five nines" IT managers have aimed to achieve for years and usually fallen short of. Now, though, the bar has been raised, and even in times of disruption, the expectation is for continuous service. The meaning of "disruption" has also been expanded to encompass security concerns: protecting data and systems from attack and compromise is a key component of ensuring they are available for enterprise use.

Technology trends within IT have made the infrastructure less diverse (by converging voice and video traffic onto data networks, for example) and at the same time more critical, as 90% of workers now reach enterprise applications provided solely from a shrinking number of corporate data centers over corporate WANs. Data thus centralized is more easily backed-up, made redundant, and controlled for compliance, so consolidating and centralizing are both driving and enabling major initiatives in all these areas.

Consolidation and centralization are also prerequisites for modern business continuity planning. Chances are, a few years ago, each data center (or server room, or server tucked in a closet somewhere) hosted something unique to it, not replicated elsewhere, and if it went down, that service was unavailable. Pulling all that function into data centers, and consolidating into the lowest number of centers that can deliver the service required (no fewer than two if extreme availability is the requirement) not only increases the need to make sure the services of the centers remain available – it also makes that possible! Once folks have found all the eggs and collected them into a basket, it is easier to replicate a matching set of eggs in a second basket, at a safe distance.

The need for continuous availability has led more than half of all participants in Nemertes' most recent data center research to use their secondary "fail over" data center not as a failover site but instead as a "hot site" providing end-user services during normal operations as well as mirroring the primary data center for continuity purposes.

Server and storage virtualization aid not only centralization and consolidation efforts but also help in enabling continuity in the face of disruptive events. Indeed, in our research we found that one of the top five business drivers for virtualization is recoverability – the ability to recover a virtual server on a different physical server in the case of hardware failure or data center disaster – or even just when it is time to patch an OS. Virtual machines are usually instantiated from an operating system and application "image" that is stored in a SAN or NAS. The recoverability advantage comes from this ability to instantiate a server by booting it directly off a SAN or NAS which can be remote from the server hardware on which the virtual machine is instantiated. Organizations can thus simplify the disaster recovery process by concentrating on the replication and availability of storage and not worrying as much about synchronization at the application layer. Storage replication at the array, SAN, or server level, or via continuous data protection (CDP) tools, can provide for real-time or near-real-time replication of data among sites to allow server images to be re-instantiated nearly instantly. Using virtualization technology in this way can be more risky to implement than simply using it for server consolidation, but the benefits can also be substantial in terms of reduced risk from downtime.

Conclusions and Recommendations

In the face of the growing emphasis on managing IT in support of business services, and the growing array of newer technologies with the potential to have transformative impact on business-line processes, enterprise IT security is and will continue to be steadily pulled towards a risk-based view of the world. Each company will therefore need to be clear on its own “risk profile,” particularly in the context of “return on risk.” How much risk, and especially, risk of loss or exposure of information, can be tolerated in exchange for what types of rewards?

With this shift, security alone is no longer a reason to not invest in key enabling technologies or practices. Instead, security staff must reposition the investments required to beef up security around such tools or practices as the cost of making these technologies (and their accompanying business benefits) possible: move from “don’t do this, it is not safe” to “here is how much we must spend to make this safe enough and available enough.” Do not say, “No sensitive data leaves the premises.” Instead, say “Laptops drives and backup tapes containing sensitive data and going offsite must be encrypted.”

Concomitantly, each organization must also reassess the relationship of the security team to IT and the rest of the organization. Splitting security out of IT, to report instead into the risk-management line or directly to top leadership, can clarify the relative roles and responsibilities and establish a more balanced corporate power structure and facilitate solid corporate governance.

About Nemertes Research: Founded in 2002, Nemertes Research specializes in analyzing the business value of emerging technologies for IT executives, vendors, and venture capitalists. Recent and upcoming research includes Web services, security, IP telephony, collaboration technologies, and bandwidth optimization. For more information about the analyst, please contact Christine Zimmerman at research@nemertes.com.