WHITE PAPER

NEXTLABS

Designing Electronic Barriers Around Digital Assets

CASE STUDIES IN COLLABORATION

INTRODUCTION

Business success today can be increasingly measured on the same scale as an organization's collaborative acumen. The best opportunities often lie beyond the traditional perimeter, in the overlap between organizations and synergies with partners, across global supply chains and a mobile workforce. However, as businesses eagerly position themselves to take advantage of these opportunities, they find themselves spending as much time—if not more—trying to understand new forms of risk.

Perhaps the biggest challenge is anticipating and preparing to block the inappropriate sharing and consumption of electronically stored information (ESI). What constitutes ***inappropriate*** sharing differs across industries, is regulated by different compliance organizations and governing bodies, and is subject to different litigation and fines. But the underlying problem is the same: ***wrongful disclosure***, where an organization fails to prevent unauthorized users from consuming certain classes of information. Common causes of wrongful disclosure include breaches in confidentiality or non-disclosure agreements (NDA) between parties, designated conflicts of interest (which are highly regulated in many industries), and failure to comply with industry and government regulations pertaining to the export of certain classes of data to foreign countries and persons. In addition to demonstrating compliance for auditing and certification purposes, businesses are expected to assume pro-active stewardship of all sensitive ESI within their organizations in the interest of investors, partners, clients, and even the public at large.

This white paper explores a technique commonly used to mitigate wrongful disclosure: implementing electronic barriers that segregate data and users. Traditional approaches to creating barriers include physically segregating infrastructure, maintaining multiple instances of applications, patching together disparate access controls and point solutions, and even "locking information up" so requests are processed on a case-by-case basis, only after labor-intensive, manual checks. These traditional approaches can be error-prone and costly to maintain, and can hamper business collaboration.

This white paper also explores a different approach to implementing electronic barriers, one that results in more precise solutions that are both data-centric and context-sensitive. Case studies of real challenges in industry collaboration reveal the need for multiple kinds of electronic barriers to target users and data in precise and complementary ways. This white paper examines the differences between three kinds of electronic barriers and explains how they can be strategically combined to prevent wrongful disclosure in four industry use cases. This paper concludes by deriving a list of the building blocks that comprise any effective electronic barrier solution.

THREE KINDS OF ELECTRONIC BARRIERS: DIFFERENCES IN DESIGN, GOAL, AND TARGETED STAGE(S) OF INFORMATION LIFECYCLE

Successful electronic barriers must be defined ***broadly*** enough to prevent wrongful disclosure, but ***precisely*** enough to allow authorized forms of collaboration to occur. Ideally, electronic barriers require neither unwieldy administrative procedures nor extensive infrastructure changes (although some procedural and infrastructure changes are often required).

Implementing successful barriers in a complex organization requires a solid understanding of ***design options*** (the shape of the logical barrier or separation, and whether it should apply to users, data, or both) and ***enforcement goals*** (the intent of enabling sharing and access or blocking sharing and access, and whether for users, data, or both). It is also necessary to understand at which point(s) in the ***information lifecycle*** electronic barriers should be applied—that is, where interventions and policy automations can be inserted into creation and storage models, access and usage patterns, modes of communication, and so on. This section examines three kinds of electronic barriers that differ in design, goal, and targeted stage(s) of the information lifecycle.

Information Barriers: Block Sharing Between Groups

Information Barriers are the simplest form of electronic barrier. They erect a logical wall between two groups of users to block communication and sharing of data. Information Barriers are commonly content-aware. That is, the restriction of information flow targets confidential information of a precise nature; the sharing of other, non-restricted information may still take place. Information Barriers are common in legal firms, where portions of an organization may serve both conflicting litigants. They are also common in the financial sector (as will be discussed in more detail in case study 1), where barriers are mandated by Federal, State, and private entities, such as the Security and Exchange Commission (SEC), National Association of Securities Dealers (NASD), Public Company Accounting Oversight Board (PCAOB), Financial Accounting Standards Board (FASB), and others.

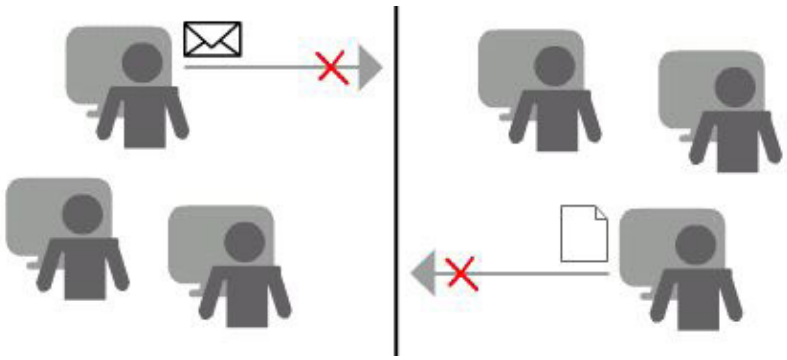

While Information Barriers typically target the sharing of a particular kind of information, they are ***user-centric*** in both design and goal. Specifically, Information Barriers are erected to prevent organizational conflicts of interest—or situations where a person or entity in a position of trust faces competing interests, sufficient to influence, or appear to influence, the objective exercise of duties. The goal of an Information Barrier is to segregate user roles so that one person may not have access to information associated with one role while simultaneously engaged in activities associated with another role deemed to be in conflict.

The most pertinent stage of the information lifecycle is the movement of data between groups. Organizations must anticipate and control modes of communication and data in transit , whether over e-mail, IM, FTP, uploading to Microsoft SharePoint or other corporate intranet, copying to file shares, publication communications, USB, and so on. Most often, Information Barriers do not exist in the ***natural infrastructure boundaries*** of an organization. Either physical infrastructure must be overhauled (which is actually required by some regulations), or policies must be implemented to control how data moves through existing infrastructure.

Information Boundaries: Restrict Sharing Within a Group

Information Boundaries encircle groups of users as the authorized consumers of confidential information, sensitive data, intellectual property, trade secrets, or any non-public information. Whereas Information Barriers are designed to block the sharing of information between users in order to segregate user roles, Information Boundaries are designed to enable the sharing of information within a group of users, but not outside of the group. Boundaries may be logical (around a facility, company, or brand). While Information Boundaries may correspond to natural infrastructure boundaries that already exist in an organization (a common example being a dedicated server or file share where a group stores data, which is only accessible to members of that group), Information Boundaries can also be more complex and unrelated to infrastructure. Think of data that must be accessed by cross-functional teams, data shared during inter-company partnerships or contract relationships, and data stored in multiple locations. Boundaries may overlap other boundaries—even entire organizational boundaries. They may also be temporary or conditional in nature, existing only for the duration of a joint venture or partnership, merger and acquisition period, or contract relationship (as will be discussed in more detail in case study 2).

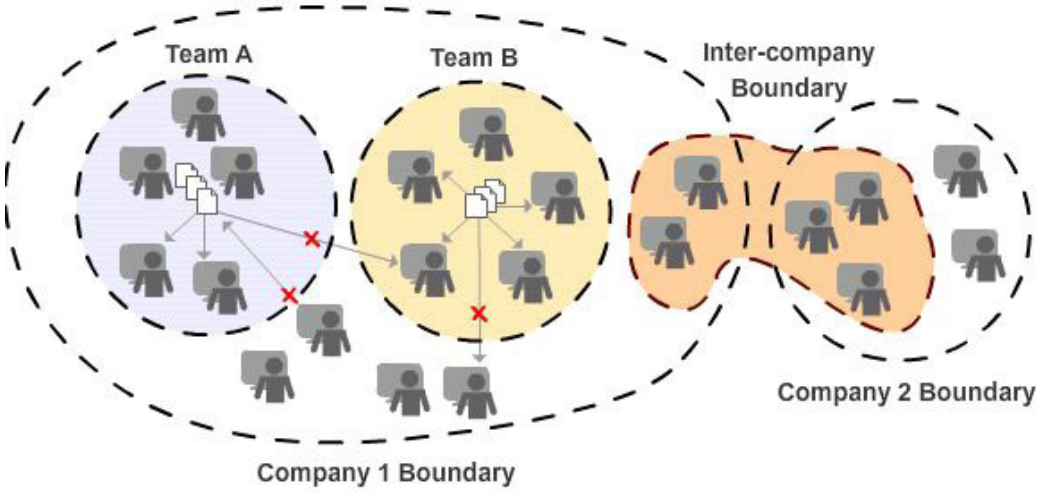

Intra- and Inter-Company Information Boundaries

While Information Boundaries can be considered ***user-centric*** in design, just like Information Barriers, they are more ***data-centric*** in goal. User groups are targeted not because there is a need to prevent conflicts of interest or segregate user roles; instead, user groups provide the easiest way to map how data should be organized into logical boundaries of access and use. In most cases, it makes sense to designate a group of engineers as the only authorized producers and consumers of data associated with an engineering project. Or, it may make sense to allow only a team associated with a particular client to have access to IP associated with that client (as is discussed in more detail in case study 3). In both cases, the point is not to block one group from sharing data with another group, but to make sure data cannot be shared outside the authorized group.

Whereas controlling the modes of communication and detecting content while in-transit are the primary lifecycle targets for Information Barriers, Information Boundaries tend to focus more broadly on usage and storage patterns. For instance, what applications are generally used to produce and consume information by users within the boundary? Where is this data stored so that others in the group may access it (in enterprise applications, in shared network locations)? How can you restrict storage and enforce access controls that reference the pertinent identity attributes (such as company, project, etc.), in order to distinguish authorized users from unauthorized ones?

Information Sandboxes: Contain Data within a Location

Information Sandboxes are a kind of boundary that have the primary goal of containing data within a logical location or system. Whereas Information Boundaries are designed primarily around ***user groups***, with the goal of restricting access and usage within that group, Information Sandboxes are designed primarily around ***data***, often corresponding to or reinforcing natural infrastructure boundaries. The primary goal of an Information Sandbox is to allow authorized users to access data only while it resides within the sandbox location.

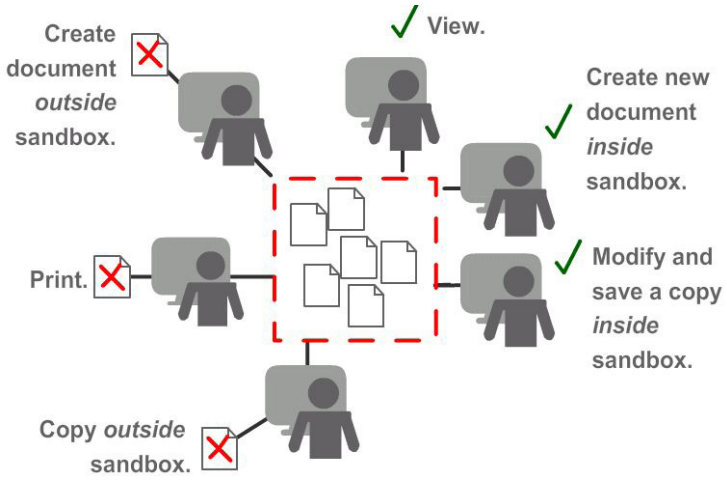

For instance, policies may specify that data of a certain class may not be created or saved outside the sandbox location. Then, additional safety controls can govern the actions authorized users can perform on data within the sandbox. While creating, viewing, uploading, and modifying may be allowed, other actions—printing, downloading, emailing, copying and pasting, moving, and renaming—may be blocked. Such controls ensure that users can perform authorized actions on digital information within the sandbox, but may not distribute or access information outside it.

Sandboxes primarily target two information lifecycle stages: blunt restrictions placed on the storage model, and fine-grained controls on access and usage patterns. In addition, as is illustrated in more detail in case study 3, sandboxes can be designed strategically to either restrict or broaden ***cross-team*** access. In other words, a sandbox may apply blunt storage rules so data may not be created or stored outside of the sandbox. Then, because data is both static and secure, limited forms of access can be extended to broader user groups, if required.

Common cases that require Information Sandboxing include Protected or Patient Health Information (PHI). In accordance with the Health Insurance Portability and Accountability Act (HIPAA), all PHI must be safeguarded against wrongful disclosure. While health care organizations should allow authorized users to create, view, modify, and store patient records in secure locations, they must also prevent authorized users from distributing those records outside secure locations. In a cases like this, organizations tend to consolidate controlled data within an enterprise application or system, and then implement access and usage controls within that application or system. In this way, Information Sandboxes often correspond with, and/or are reinforced by, natural application and infrastructure boundaries.

The need to comply with export regulations governing controlled technical data is another common industry driver for Information Sandboxing. In accordance with regulations such as the International Trade in Arms Regulations (ITAR) (to be discussed in more detail in case study 4), controlled technical data should only be stored and routed on IT infrastructure located on US soil.

While the benefit of Information Sandboxing is the ability to contain data within a secure location, sandboxes often need to be combined with other kinds of barriers to enable successful collaboration. An overly broad sandboxing approach, in other words, may impede necessary sharing of data.

	Design	Goal	Information Lifecycle	Classification Requirements
Information Barrier	Simple barrier between groups Usually content aware Usually not reinforced by natural infrastructure boundaries, although infrastructure changes are sometimes mandated by regulations	To prevent sharing between users To segregate roles to prevent conflicts of interest	Ability to anticipate and control all modes of communication Ability to detect restricted data while in-transit (sent via email, uploaded to file shares, SharePoint, etc.)	*Ability to detect and classify restricted content, either by content or label, in the midst of communication or sharing* *Ability to detect identity traits for information senders and recipients*
Information Boundary	Boundary encircling user groups Often multiple, overlapping, and temporary Can correspond with logical architecture, system, application and physical location, or not	To enable access, usage, and sharing of data by all members of a user group	Knowledge of core applications used to create and consume data, ability to enforce access controls within these applications Knowledge of and ability to control the storage model for controlled data	*Ability to apply classifications to data at rest and during use (at the point of initial creation, storage, and access)* *Ability to distinguish authorized users from unauthorized users, based on pertinent identity attributes*
Information Sandbox	Boundary containing data Designed to be permanent Often reinforced through logical architecture, system, enterprise application, and/or physical location	To enable authorized access and use while data resides within a logical, physical, or systemic location; to contain a certain class of data	Ability to restrict storage, creation, and usage of data within a logical boundary Granular control of user behaviors to permit certain actions and block others within the sandbox	*Ability to classify data that requires sandboxing persistently across all stages of the information lifecycle* *Ability to distinguish authorized from unauthorized users based on pertinent identity attributes*

CASE STUDIES FROM INDUSTRY

Most often, policy designers should strategically combine electronic barriers (Information Barriers, Boundaries, and Sandboxes) to create precise solutions that effectively mitigate the risk of wrongful disclosure. Consider the following industry case studies.

Case Study 1: Conflicts of Interest in the Financial Sector

Information Barriers have received the most media attention in the Financial Industry, where high profile legal cases and legislation have resulted in broad reforms to regulate conflicts of interest.

In one of the most highly publicized settlements, in a case brought by the SEC, NASD, and NYSE (often referred to as the "Global Analyst Research Settlement"), ten of the largest Wall Street firms agreed to pay over one billion dollars in fines for publishing stock research that was manipulated to benefit firms' and clients' investment banking activities. Around the same time, as a reaction to public outcry over major corporate scandals involving ENRON, WorldCom, and others, Congress passed the Sarbanes-Oxley act, also known as the "Public Company Accounting Reform and Investor Protection Act." This legislation expanded and reinforced reforms called for in the Global Analyst Research Settlement, including requirements to establish oversight boards, ensure auditor independence, and prevent "Securities Analyst Conflict of Interest."

One outcome of these events is that financial organizations are now tasked with blocking all communication between the wing of their organization that provides loans and handles mergers and acquisitions, and the wing that produces securities analysis and research to make buy and sell recommendations on stocks and bonds. The precise conflict of interest to be prevented is where a research analyst's objectivity may be tainted by knowledge of ongoing investment banking initiatives. The temptation may exist to make erroneously favorable or unfavorable projections in service of investment banking activities the analyst is privy to. To prevent this cross-contamination of user roles, financial organizations are tasked with physically separating departments and preventing the sharing of information between them.

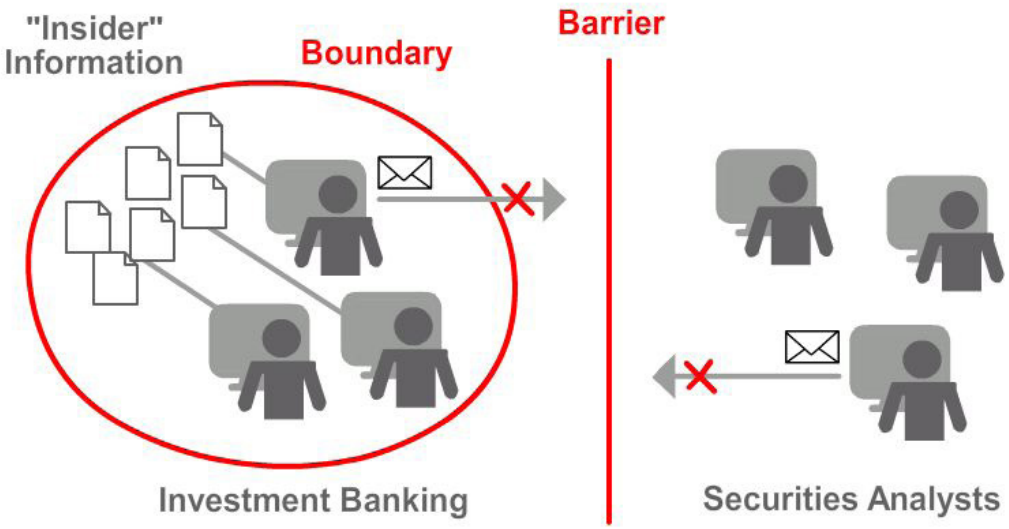

To satisfy these regulations, avoid costly legal battles and fines, and reestablish public and investor confidence, financial institutions should prevent wrongful disclosure using two kinds of electronic barriers: Information Barriers to block the communication and sharing of certain kinds of information between investment banking and securities analysts, and Information Boundaries around Investment Bankers who have access to "insider" information (non-public, unpublished research about the health of financial institutions). Investment bankers should be able to create, access, store, or share this information with each other in authorized ways; however, this information should be inaccessible to securities research analysts, as well as anyone outside investment banking. This solution requires the ability to detect "insider" information, control how it is stored and accessed, apply communications controls that can detect insider information while in transit, and block communications based on the identity of senders and recipients.

Case Study 2: Intellectual Property in Joint Ventures

Another challenge organizations face is protecting intellectual property (IP) in the midst of, or following, a period of inter-company collaboration. Consider the case of joint ventures, common in government projects, Aerospace and Defense (A&D), construction, and other large scale efforts. Joint ventures are common in cases where the technology is multi-faceted and complex, thus requiring collaboration among entities with highly specialized areas of expertise. Because joint ventures allow multiple entities to pool resources and share labor, they empower businesses to accomplish together what they could not accomplish alone. Plus, joint ventures and partnerships can be an especially vital source of entrepreneurial synergy and technical innovation in the formative stages of a business, technology, or industry.

On the other hand, many well-known cases attest to the hazards posed from sharing IP with partners, since those partners tend to become eventual competitors. One well-known example is the case brought by Lexar against Toshiba. In 1997, California-based Lexar and Japan-based Toshiba entered into an agreement to share IP while co-developing flash memory technology (the technology behind digital photography, consumer electronics, and other industrial and communications technologies). In a 2005 court case, Lexar sued Toshiba for leaking vital trade secrets to Lexar's main competitor, Sandisk, with whom Toshiba had formed a subsequent joint venture. The court awarded Lexar a $464 million dollar settlement.

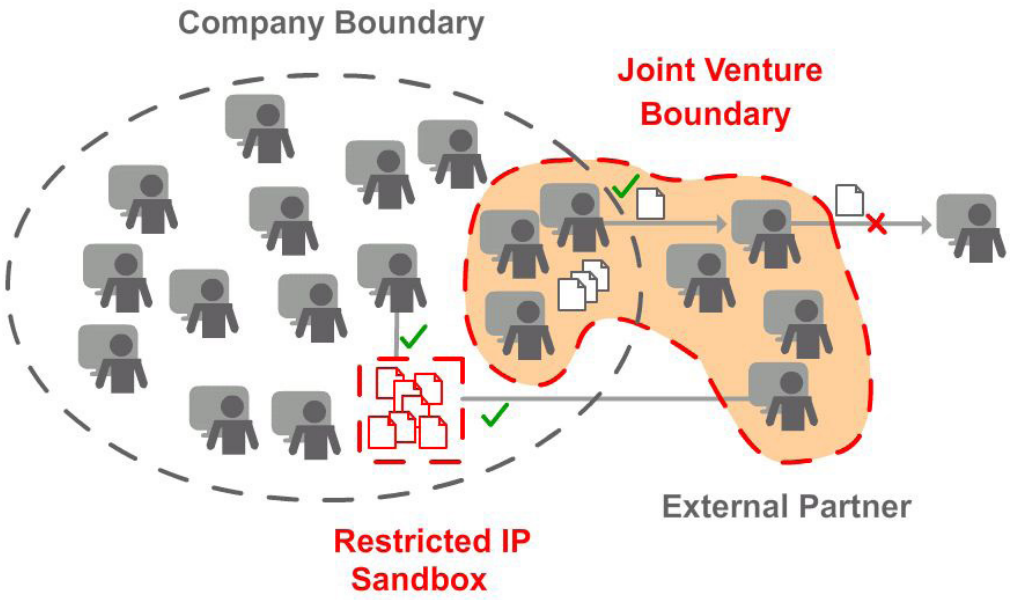

To mitigate the risk associated with a case like this one, organizations should apply precise and persistent controls to govern how their IP is shared with partners. First, an inter-company Information Boundary should be drawn around users associated with the joint venture. Sharing of IP would be permitted within this team, but prevented outside of it. When the period of the joint venture is complete, the boundary should dissolve and data should no longer be accessible by external partners. The joint venture boundary thus requires controls to be enforced on devices off the main corporate network, as well as the ability to rescind access and usage rights when the partnership is complete.

Next, more highly sensitive IP could be sandboxed so that it can be controlled even more strictly. By limiting access to the most sensitive or critical IP, a sandbox would allow authorized users to perform actions on data only while it resides on the controlled system or location. The Sandbox could also apply granular controls to allow authorized users to enter the sandbox and work with data in appropriate ways, but block users from downloading to an endpoint, sharing with other users, and so on.

Importantly, this combined Information Boundary and Sandbox solution would enforce distinct storage and usage patterns for different classes of data: one class of IP is approved for access, download to an endpoint, and distribution to partners' computers, while another more highly restricted class requires sandboxing. The solution thus requires classifying data and users, and user classifications may need to encompass multiple identity attributes at once: company employees associated with the joint venture, external partners associated with the joint venture, and various internal employees granted access to the sandbox.

Case Study 3: Customer Intellectual Property in Contract Manufacturing

Another kind of wrongful disclosure, common in manufacturing services and contract manufacturing, stems not so much from the need to protect an organization's own IP during or after a period of collaboration (as with case study 2), as from the need to segregate and protect the IP of one or more clients. For example, it is not uncommon for a single contract manufacturer to be working with IP from multiple clients who are in direct competition with one another. Increasingly, to win these contracts, manufacturers must demonstrate their ability to safeguard and segregate all client IP. Oftentimes, this means conducting regular auditing and reporting to demonstrate the effectiveness of procedures and employee compliance.

Information Barriers, Boundaries, and Sandboxes could be combined to effectively protect IP in this scenario. For example, an Information Boundary could ensure that only Team A can access and use IP associated with Client A, the same for Team B and Client B. An Information Barrier could then also be applied to prevent the communication between teams working on competing clients' IP. In this case, users can download client IP and email it to other members within the authorized group. But, a content-aware Information Barrier would block Team A from sending emails containing Client A information to Team B (whereas Team A could share information regarding other clients or unclassified data, such as Team U in the figure).

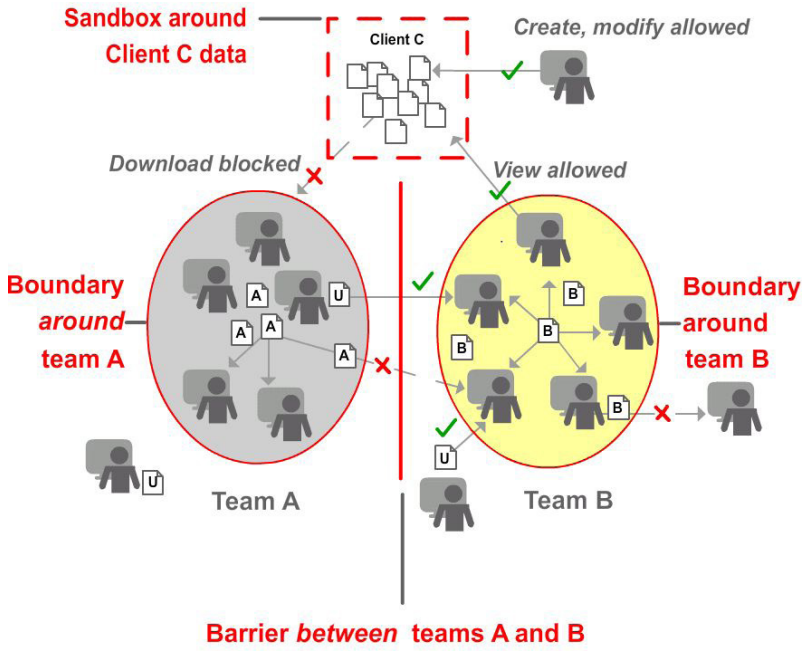

Finally, an Information Sandbox could be added strategically, to address two different use cases. One use case is where another client has more stringent IP protection requirements than others, so their data must be stored within an Information Sandbox. In this case, client IP can only be accessed and used while it resides on a secure server or within the sandbox application. A second use case (illustrated in the image above), would be a sandboxing approach that is designed to enable broader cross-team access to a client's data. It may be the case, for instance, that inadvertent disclosure of IP is the main concern, ***not*** the contamination of user roles. Sandboxing a client's IP could enable users from different client teams to safely access and handle client C's data. Because data will never leave the secure location, the threat that users will inadvertently disclose data (in an email sent to a competitor, for instance) is diminished. Note how this last use case illustrates the different goals of designing boundaries around ***user groups*** versus ***data***.

Case Study 4: Export Compliance

Many businesses are grappling with the complexity of export compliance regulations for controlled technical data. Taking the International Traffic in Arms Regulation (ITAR) as just one example, regulations tend to prohibit the export of "Defense Articles," which typically refers to products as well as technical data and services associated with those products. Export to non-US countries without a proper export license constitutes a violation, as well as sharing or giving access to foreign persons, including access by foreign workers on US soil. Routing or storing technical data through servers or storing data in file shares located in other countries can also constitute a violation.

Clearly, these regulations present hurdles to businesses that seek to collaborate and share export-controlled technical data across global supply chains, in international trade with overseas customers, or while manufacturing for foreign clients. Several high profile cases attest to the prevalence of this business challenge, including:

- Boeing (2009), fined $15 million dollars for violations of the ITAR; found guilty of ITAR Compliance violation every other year since 9/11
- Lockheed Martin (2008), fined $4 million dollars for exporting technical data connected to the sale of Hellfire missiles to the United Arab Emigrates
- Northrop Grumman (2008), fined $15 million dollars for export of controlled parts and technology for commercial inertial navigation units
- ITT Corp (2007), fined $100 million dollars for exporting military night vision goggles to foreign countries
- Loral Space & Communications Corp., Boeing, and Hughes Electronics (2003), fined $14 million dollars for exporting satellite launch rocket integration and failure analysis services to China; while the satellites were commercial, the launch vehicles were subject to ITAR

Many of these high profile violations occurred because sub-components and processes associated with a larger assembly are subject to ITAR. Smaller defense articles may "contaminate" otherwise uncontrolled commercial products. The inclusion of an ITAR component, no matter how small, will cause a larger commercial assembly to be subject to ITAR. Its export, the export of any information about it, the storing or routing of information about it on servers on foreign soil, and the sharing of information about it with a non-US citizen would all constitute a violation.

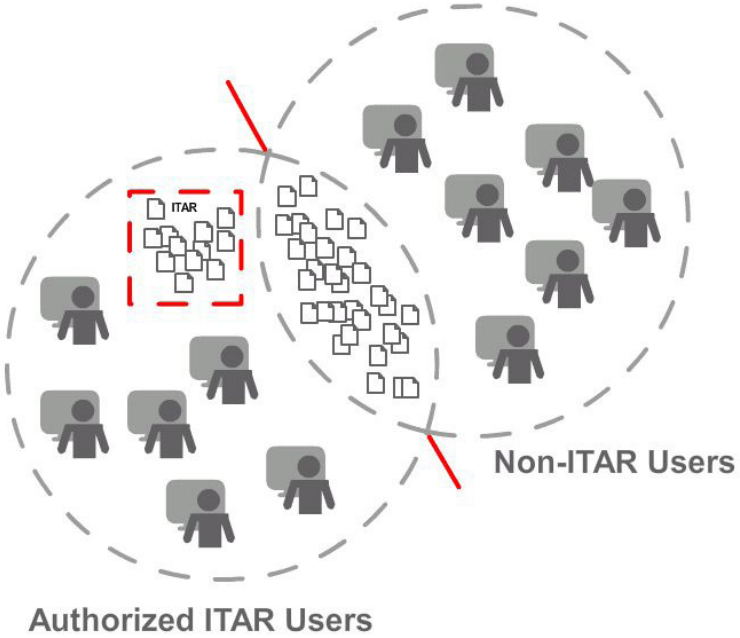

Organizations trying to pursue business opportunities associated with controlled technical data, whether for global manufacturing, international sales, or military contracts, tend to have users from multiple countries and data located in multiple locations. In addition, business will often have data that ***is*** and ***is not*** subject to compliance regulations. In fact, for some cases, organizations are dealing with a small subset of their business that is subject to export compliance regulations, but because of the "contamination" issue and the fact that they have employees in multiple locations and of varying citizenship, they must properly segregate this small subset of data before conducting other business. For complex cases like these, Information Barriers, Boundaries, and Sandboxes may all be necessary.

First, all data subject to ITAR, for which no export license exists, would need to be sandboxed within logical and physical locations within US soil. This addresses the ITAR requirement that controlled technical data never be stored or routed through a foreign country without the proper export license. Second, Information Boundaries would need to be applied to limit access to ITAR data to ITAR-authorized users. Creating this Information Boundary can be more difficult than it first sounds, as it must combine identity attributes for citizenship, current geographical location, and pertinent export licenses. In this case, authorized users would include US employees in US locations and users in other locations with export licenses. Unauthorized users would include US employees in non-US locations, non-US employees in US locations, and non-US employees in non-US locations (or more simply, those "not in" the authorized user group).

Use Case	Solution Components	Required Solution Capabilities
1. Preventing conflicts of interest in the financial sector	Information Barriers across modes of communication and collaboration Information Boundary around research analysts to govern access, storage, and usage models	Classification of "Insider" information Identification of users based on identity attributes (department) Detection of classified data while in-transit Access and usage controls for all applications used to create and modify data Location-based controls to govern where data is stored
2. Protecting intellectual property during and after joint ventures	Information Boundary around users associated with the joint venture, to restrict communication, collaboration, access, storage, and usage Possible Information Sandbox to contain highly critical IP	Classification of semi-critical and critical IP Identification of users based on identity attributes (joint venture team, authorized internal users) Access, usage, and storage controls for all applications used by authorized users to work with semi-critical IP Off-the-network digital rights enforced for semi-critical IP that is downloaded to workstations outside the organization Access, storage, and usage controls for data residing within the Information Sandbox
3. Segregating client intellectual property	Information Barriers and Boundaries separating client teams when there are concerns about segregation of duties, controlling modes of communication, access, usage, and storage Information Sandbox for client IP when added protection is required or to enable safer cross-team access	Classification of semi-critical and critical IP Identification of users based on identity attributes (joint venture team members, authorized internal users) Access, usage, and location-based storage controls for all applications used to handle Client A and B data Access, usage, and storage controls for application or system used to contain Client C data
4. Complying with export control regulations	Overlapping Information Boundaries allowing US and licensed users to access ITAR data Additional Information Boundaries for users/locations where there are proper exceptions/licenses Information Sandbox restricting creation and storage of classified data to secure US locations	Classification of ITAR restricted data Identification of users based on multiple identity attributes (citizenship, location, applicable export licenses) Location-based access, usage, and storage controls for all applications used to handle ITAR data within the Information Boundary Location-based access, usage, and storage controls for applications or systems used to sandbox ITAR data

Once these groups are mapped out, overlapping Information Boundaries could be drawn around the groups so ITAR-authorized users can access both non-ITAR and ITAR data, whereas non-ITAR users can access only non-ITAR data (note that the example in the image does not factor in the complexity of export licenses). Finally, a content-aware Information Barrier could be erected between authorized and unauthorized users to block the sharing of any ITAR data between the two groups. This could be necessary for ITAR data that is exported to a foreign location where a license exists, but where there is no Information Sandbox to lock data in place.

SOLUTION BUILDING BLOCKS FOR IMPLEMENTING DIGITAL INFORMATION BARRIERS

The industry case studies discussed above illustrate how effective electronic barriers—that is, those capable of blocking wrongful disclosure while enabling appropriate forms of collaboration—must be precise and granular enough to be designed around data and/or users to accomplish specific, limited objectives. While the solution for each case study differs, together they reveal the basic building blocks of any good electronic barrier solution:

- Discovery, Analysis, and Persistent Classification of Data
- Attribute-Based Classification of Users
- Flexible Location-Based Controls
- Centralized, Standards-Based Policy Engine with Electronic Barriers Support
- Easy-to-Define Access, End-Use, and Storage Controls
- Enforcement across Multiple Locations, Applications, and Communication Channels
- Persistent Data-Level Controls off the Server or Network
- Monitor and Audit of Electronic Barriers

Discovery, Analysis, and Persistent Classification of Data

All four solutions discussed in the case studies above require some means to classify data—in case 1, to designate "insider" information, in case 2, to distinguish between sharable IP and restricted IP that must be sandboxed, in case 3, to label customer IP, and in case 4, to identify data subject to ITAR export compliance regulations and data for which an export license has been granted.

Classifying existing data usually entails a ***discovery*** process for all data stores, servers, databases, applications, and even user workstations. Organizations often handle this through a batch scanning process that applies classifications to existing data based on the results of ***content analysis***, which may consider file attributes (owner, creator, date created), content (key words or types), file type, and/or location. In addition to ***classifying*** data that already exists, organizations may classify data at the point of creation or while data is in transit. This can happen through an automated content analysis and labeling process that "kicks off" whenever users perform an action (such as save a file, upload to a file share, etc.). Alternatively, classification can be a user-driven process where, based on the presence of certain key words, the identity of the creator, and/or the origination or storage location, users are prompted to supply the appropriate classification themselves.

Importantly, once data is classified, these classifications should be ***persistent*** if data is permitted to move from one location to another. Take a scenario where sensitive client IP is stored in a designated Microsoft SharePoint site or other corporate intranet, which is only accessible to a team of users associated with that client (similar to the Information Boundary discussed in case study 3). The classification of that client IP must persist when the data is downloaded to authorized users' desktops. Without persistence, a "content aware" Information Barrier could not be enforced, that is, a barrier that blocks the emailing of classified content between groups of users, but enables the sharing of other non-sensitive content.

Attribute-Based Classification of Users

All four case studies discussed above also require the ability to classify users based on their characterisitcs, or ***attributes***. More complex cases often require the ability to classify users based on ***multiple attributes at once*** (citizenship, location, department, project, and so on), and may also require attributes that are housed in different identity stores. For case study 3, in order to determine whether users are authorized, controls must be able to detect whether they are internal employees or external partners, and for internal employees, whether they are members of the joint venture team or not. In this case, identity attributes of users outside the organization must be collected from an external identity store, such as the Active Directory of the partnering organization, and then be integrated and applied in controls alongside the identity attributes of internal users.

Another example is case study 4, where maintaining Information Boundaries for ITAR compliance requires the classification of users based on citizenship, the presence of export licenses, the location of access, and possible other business attributes, such as team membership or job title. While some of these attributes are likely static or semi-static (as in citizenship), others are likely to change or be time-sensitive (as in location and possible team membership). Some of these traits may be housed in a separate application outside Active Directory—for example, identity attributes assigned to users in a shared Human Resource Management System (HRMS), Enterprise Resource Planning (ERP), or Product Lifecycle Management (PLM) application. An effective Electronic Barrier must be capable of referencing these identity stores to evaluate user attributes dynamically at the point of policy enforcement.

Flexible Location-Based Controls

An effective electronic barrier solution must also be able to enforce controls based on the locations of data and users. Location-based controls should be flexible enough to either lock data in place (as in a sandbox location), or apply controls dynamically to data as it moves across multiple locations. From within the same policy set, a business may need to enforce certain controls while data resides within file servers and shared network locations, other controls when the same data is downloaded to a user's desktop, and other controls when that data is emailed to an external user outside the organization. The same is true for the location of users: one control may need apply to a group of users whenever they access data remotely, another control when users travel outside the country.

This requirement is apparent in all four use cases, but is particularly clear in use case 2, where an organization wants to share data with a partner during a designated project period. Highly critical IP may need to be sandboxed within a dedicated file share location. Other less critical IP may reside within more public company file shares, where authorized users can view, download, and print documents. When data is downloaded to an authorized user's endpoint, other controls allow users to modify documents, save versions or copies, upload documents to file shares (but not overwrite originals), and email documents to authorized users in the external partner organization. Finally, when data is within the possession of that partner organization, additional controls may need to restrict redistribution and modification. In all cases, the solution should be able to dynamically detect ***the location*** from which an access or usage request originates, and enforce the appropriate control.

Centralized, Standards-Based Policy Engine, with Electronic Barriers Support

To sum up the building blocks discussed thus far, an effective electronic barrier solution should be able to: (1) apply data classifications, (2) reference and combine disparate identity attributes, and (3) enforce controls dynamically based on the locations of data and users.

These requirements point to the need for a ***centralized***, ***standards-based*** policy management tool. A ***centralized*** tool is necessary so all data classifications, identity attributes, and location definitions can be gathered and managed within ***one application***, by ***one team***, in the ***same set of policies***. The ***standards*** requirement ensures that the policy authorization system will be both extensible and compatible with other authorization tools (the current industry standard for policy languages being eXtensible Access Control Markup Language (XACML)).

The authorization mechanism should be externalized from the application logic of the applications where controls are applied. In other words, they should apply to more than just Microsoft SharePoint, Adobe Acrobat Reader, or Microsoft Office files. At the same time, the authorization mechanism must be flexible enough to integrate or "hook into" all these applications. Ideally, the policy platform should also include built-in support for the kinds of controls that are typically needed for effective electronic barriers, namely: endpoint controls, communication controls, access controls for data on servers, and pertinent collaborative and enterprise application controls.

Centralized, Standards-Based Policy Engine, with Electronic Barriers Support

To sum up the building blocks discussed thus far, an effective electronic barrier solution should be able to: (1) apply data classifications, (2) reference and combine disparate identity attributes, and (3) enforce controls dynamically based on the locations of data and users.

These requirements point to the need for a ***centralized***, ***standards-based*** policy management tool. A ***centralized*** tool is necessary so all data classifications, identity attributes, and location definitions can be gathered and managed within ***one application***, by ***one team***, in the ***same set of policies***. The ***standards*** requirement ensures that the policy authorization system will be both extensible and compatible with other authorization tools (the current industry standard for policy languages being eXtensible Access Control Markup Language (XACML)).

The authorization mechanism should be externalized from the application logic of the applications where controls are applied. In other words, they should apply to more than just Microsoft SharePoint, Adobe Acrobat Reader, or Microsoft Office files. At the same time, the authorization mechanism must be flexible enough to integrate or "hook into" all these applications. Ideally, the policy platform should also include built-in support for the kinds of controls that are typically needed for effective electronic barriers, namely: endpoint controls, communication controls, access controls for data on servers, and pertinent collaborative and enterprise application controls.

Easy-to-Define Access, End-Use, and Storage Controls

The policy engine should include an easy way to define all the various actions users can perform in the multiple applications where controls are applied. Another way to phrase this requirement is that any effective electronic barrier solution should be ***business-aligned***, that is, capable of allowing or denying user requests for the typical work of a regular business day. Taking the example of actions that typically happen at a user workstation, this list might include viewing, modifying, printing, copying, emailing, performing a save as, and saving to removable media, to even more granular actions, such as taking a screen capture, copying content from a file to the clipboard, and changing file properties. Beyond being able to apply basic Windows desktop controls at this level of granularity, the solution should also support controls for access on servers, uploading to web portals, and actions performed in specialized collaborative applications.

The need for controls that can be defined at this level of granularity is obvious in the case of Information Sandboxing. While the data location remains static, permissions to be granted to that data may be highly granular, so that the data can be created and viewed, but not moved or distributed outside the sandbox location. In case study 3, it would be necessary to prevent an employee working with the IP of client A to copy data from a sandboxed application into another document intended for circulation with Team B. Information Barriers and Boundaries also require precision and granularity in how user controls are defined. For example, it may be necessary to prevent users from uploading data to public file share locations accessible by users associated with other departments (as in case studies 1 and 3).

Enforcement across Multiple Locations, Applications, and Communication Channels

The case studies discussed above do not specify exactly where data is being stored (in file shares, servers, endpoints, web portals and Microsoft SharePoint, or in which collaborative enterprise applications (such as ERP, PLM, and Document or Content Management Systems (DMS, CMS)). The case studies also do not specify which communication channels are being used to share data (email, instant message, web meeting sessions, etc.). The assumption is that the specifics will change from business to business, and that a successful electronic barrier solution will need to apply controls across all the pertinent locations, applications, and communication channels.

For example, it may be necessary to control access to ITAR technical data across several locations at once, as with case study 4: data posted to Microsoft SharePoint, stored on file shares, and inserted into SAP, a common ERP application. Or, in a scenarios like case study 1, classified "insider" data should be stored only in certain locations (a file share where only authorized users can upload, view, and download files). Authorized users may download this data to their endpoints, but they are prevented from emailing data to certain users.

For both cases, a successful solution must apply consistent controls around data as it travels through locations and across communication channels. Organizations soon discover that enforcing controls within only one location, application, or channel is much simpler than applying controls consistently as data moves from one location to the next. To respond to this requirement, organizations often try to implement disparate point-specific solutions, which often results in extra administrative work, hampered collaboration, and the inadvertent blocking of authorized access and sharing requests.

Persistent Data-Level Controls, Off the Server or Network

The ability to control how documents are accessed and used while they reside in controlled Information Sandboxes or on the corporate network is one thing; the ability to apply controls off the server or network is another.

This requirement is illustrated in detail in case study 2, where an organization needs to share data outside the traditional company perimeter in order to enable collaboration with outside contractors, suppliers, and partners. Oftentimes, this data should be shared with users conditionally—only for the duration of a contract or joint venture. In addition, an organization may want to apply strict controls to what users can do with this data while it is outside the organization (permitting downloaded and viewing, but not emailing, for instance).

These kinds of scenarios require an electronic barrier solution that can apply persistent data-level access and usage digital rights, even on devices outside the organization's IT infrastructure. While policies should ideally be managed and distributed by a central policy server, enforcement should not be dependent on constant server connectivity.

Monitor and Audit of Electronic Barriers

A final piece of an effective electronic barrier solution is the capability to monitor activities, conduct audits, and report on events. This requirement is easily visible in use cases 1, 3, and 4, where organizations must run regular reports to demonstrate compliance with financial industry regulations, win contracts when clients request proof that their IP will be secure, and avoid export compliance fines. An effective solution should be able to leverage logs that capture user activity across the enterprise. Reporting tools should also be flexible enough to satisfy multiple auditing and reporting requirements.

Audit, monitoring, and reporting features can also be important in early planning and implementation stages, when an organization is trying to understand the best way to design effective electronic barriers. For instance, "monitor only" policies may be deployed to audit how data is accessed and used on file servers, or how it is uploaded and downloaded from Microsoft SharePoint sites. Reports can include which users are downloading what data, to what locations, so that businesses can better understand storage, access, and usage patterns and what kinds of controls are needed.

ABOUT THE NEXTLABS SOLUTION FOR INFORMATION RISK MANAGEMENT

NextLabs' approach to electronic barriers is designed with today's collaborative challenges in mind. The NextLabs solution helps companies comply with industry regulations by enforcing and auditing barriers during all stages of the information lifecycle. Using NextLabs' centralized policy management system, organizations can combine controls to design and implement Information Barriers, Boundaries, and Sandboxes to meets their particular industry and business requirements.

The policy engine is powered by a standards-based XACML policy language with built-in electronic barriers support. The XACML-based policy engine allows easy creation of controls, in the form of policies or rules, to enable combinations of various types of electronic barriers. A complex electronic barrier requirement can be implemented using a single set of policies, with fine-grained controls across multiple teams, applications, and modes of communication. Both business analysts and IT administrators will be able to create policy using common business terms and vocabulary to reflect underlying business processes, workflows, and compliance requirements. This allows for the creation of a common policy language and fabric for IT and business leaders to collaborate and address complex compliance and risk management objectives. As a result, controls can be integrated with critical business applications and applied uniformly to ensure consistent enforcement and to create a cohesive solution deployed across the enterprise. Administration and management of the policies can be delegated to the corresponding business stakeholder, as well as the Policy Administrator. Additionally, policy objects and audit data can be managed centrally with web-based administration and reporting tools.

Evaluation of policies is performed dynamically, in real-time, based on attributes of the information, user, event, and environment. Information risk management or data security policies, including electronic barrier policies, can be identity-driven and content- and location-aware, while also supporting a spectrum of enforcement options: allowing or blocking user actions, educating users on proper data handling procedures and applicable policies, and automating workflows and providing other remediations. NextLabs' central policy management engine can also leverage identity attributes from multiple data stores, for fine-grained Attribute Based Access Control (ABAC), with out-of-the-box support for endpoints, communications, servers, Microsoft SharePoint, SAP, PLM applications, and more.

Specifically, the NextLabs' solution allows companies to:

- Create Information Barriers, Boundaries, and Sandboxes to reflect business relationships and use cases typical of IP protection, financial regulation, and export compliance
- Manage data access, handling, and disclosure with consistency across communication and collaboration channels to prevent improper activities, while not interfering with "business as usual"
- Centralize policy and classification management in a XACML-based policy server, with support for data discovery, content analysis, and persistent classification, as well as dynamic, extensible integration with multiple identity stores
- Monitor activities comprehensively (surveillance), simplify auditing, and report violations to prove effective policy; Educate users about policies and procedures to increase compliance awareness
- Apply one set of common policies across and outside the enterprise: workstation, server access control, common enterprise collaborative applications (Microsoft SharePoint, SAP ECC, SAP PLM)

In addition, NextLabs provides out-of-the-box Electronic Barrier support, with pre-built policy objects and components designers can use to construct policies. Policy sets are interoperable and easily customized to the environment to address areas of risk, including the following areas:

Unified Communications

Unified Communications policies can be applied across multi-channel communications to apply consistent controls across voice and electronic communications applications (IM, e-mail, VoIP, Web conference, etc.). These policies are integrated with strong automation and remediation capabilities. Example policies include:

- When a chat is initiated over instant messenger between users with a potential conflict of interest, automatically add a chaperone to monitor the conversation.
- Deny employees on a Web conference located outside of the region from saving client data distributed electronically.
- When the EU branch office attempts to email client account information outside the region, quarantine documents and initiate approval procedures.

Collaboration

Collaboration policies apply controls across collaboration tools, such as Microsoft® Office Microsoft SharePoint® portals, or SAP ECC or PLM. Example policies include:

- Prevent anyone outside of the research team from accessing unpublished research in designated Microsoft SharePoint document libraries (regardless of access rights delegated by Microsoft SharePoint administrators).
- When non-EU employees attempt to access and download EU client account files, warn the employees about regional regulations and log the attempt for auditing.
- When a non-US user attempts to access an ITAR classified material in SAP ECC or PLM, check for the presence of an export license, and when lacking, deny the access request; log the attempt for auditing and display an educational user alert.

File Sharing

File Sharing policies apply controls across desktops, Microsoft Windows® and Linux® file shares, and Web or FTP servers to limit disclosure. Example policies include:

- Allow account managers of the company's Japan subsidiary to upload client account records only to Japan regional servers.
- Prevent client team "A" from accessing Merger and Acquisition (M&A) deal files stored in the Windows file share directory used by client team "B" who is responsible for a competing client.

ABOUT NEXTLABS

NextLabs®, Inc. provides data-centric security software to protect business critical data and applications. Our patented dynamic authorization technology and industry leading attribute-based policy platform helps enterprises identify and protect sensitive data, monitor and control access to the data, and prevent regulatory violations – whether in the cloud or on premises. The software automates enforcement of security controls and compliance policies to enable secure information sharing across the extended enterprise. NextLabs has some of the largest global enterprises as customers and has strategic relationships with industry leaders such as SAP, Siemens, Microsoft, and IBM. For more information on NextLabs, please visit http://www.nextlabs.com.

NEXTLABS