

NEXTLABS

Implementing Data Security Using Attribute Based Access Control (ABAC)

The NextLabs Solution Architecture



Attribute Based Access Control (ABAC) has proven to be the best approach to data-centric security to keep pace with the demands of today's extended enterprise. Organizations can leverage current business processes and policies to implement ABAC and automate their security enforcement.

This white paper explains how NextLabs' solution enables real-time authorization and dynamically evaluates information access events by using the most up-to-date information. Security Professionals, IT Architects, and System Integrators will understand the benefits of implementing data-centric ABAC and the best practices to shorten time to market.

INTRODUCTION

The Challenge of the Extended Enterprise

Security professionals already grapple with uncoordinated information infrastructure and a patchwork of disparate security systems. But now the “extended enterprise”— what Forrester describes as an “ecosystem of customers, devices, clouds, service providers, partners, supply chains, and empowered users”—is highlighting the fundamental weaknesses of traditional identity and access management (IAM). Organizations must share critical business information with partners, customers, and other third parties to remain competitive.

- Data is not in “containers” or “locked up” – it is much more accessible to employees and partners
- It is difficult to identify all authorized users and devices
- User and partner status change frequently and static systems cannot keep up

The ABAC Solution

Organizations implement Attribute Based Access Control (ABAC) because they acknowledge traditional IAM is not adequate for the challenges of the extended enterprise, including how to safely share confidential information and adhere to regulatory data requirements.

ABAC allows you to design controls around the characteristics of data that warrant protection in the first place; this could be type of content, project, security clearance, and so on. Attribute Based Access Control also takes into account information about the user and the environment, including location, position, device, and network.

Controls can be written as simple versions of information sharing policies. Once written, a single policy can be deployed across multiple systems and hundreds of devices. Unlike traditional controls, which require permissions to be defined statically before an access attempt occurs, ABAC rules are evaluated dynamically with attributes presented at run-time. The attributes can come from multiple sources – even sources external to an organization.

Plus, enforcement adapts to risk level automatically. For example, if the classification of a document changes, or a user’s team membership changes, access rights are automatically adjusted. No need to request new roles or update permissions.

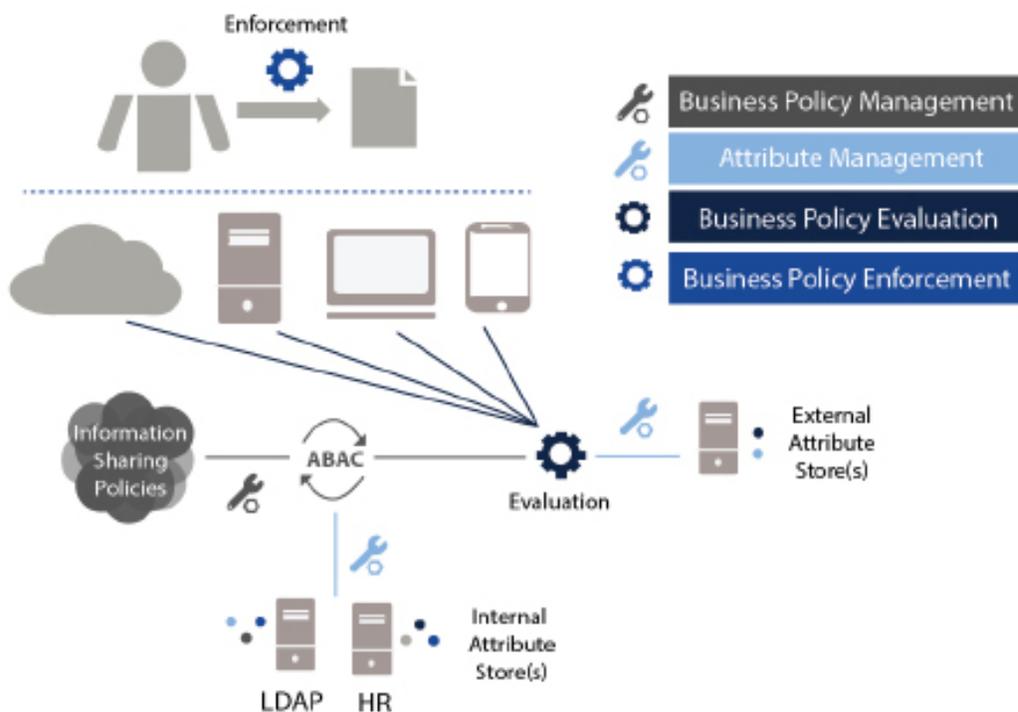
Implementing ABAC

ABAC requires various inputs to work. For data-centric use cases, we need three critical inputs: data classification, identity attributes, and policies. Most organizations have the processes in place to provide these inputs. NextLabs Control Center platform was designed to automate these processes and enable data-centric ABAC policies to be centrally applied across the extended enterprise.

THE NEXTLABS APPROACH

NextLabs' Control Center provides services, integration points, and automation tools to allow organizations to centrally administer, deploy, and enforce data-centric ABAC policies. The tools can be grouped into the following Control Center components:

- **Business Policy Management** – Digitize information sharing requirements as policy, centrally manage and deploy policies
- **Attribute Management** – Leverage existing attributes, delegate ownership, and define integration points to internal and external attribute stores
- **Business Policy Evaluation** – Dynamically evaluate data access no matter where data resides, using attributes presented at run-time
- **User and Data-Centric Enforcement** – Automate user and data-centric information controls



Business Policy Management

With NextLabs, business policies are digital versions of your information sharing requirements. Business policies are centrally managed in natural language and deployed cross-system. Business Policy Management enables organizations to apply one set of business policies across applications and systems, rather than “translate” information sharing requirements multiple times into permissions, roles, and ACLs.

Business Policy Management employs the following Control Center technologies (discussed in more detail below):

- **Policy Language** – Any user can create policies in business or natural language
- **Policy Component Model** – Policy building blocks reflect local organization structure.
- **Policy Lifecycle Management** – Easy drag-and-drop user interface, so that users create and manage policies.

Active Control Policy Language

Control Center’s Active Control Policy Language (ACPL), supports fine-grained policy for enforcing information sharing requirements. ACPL is a fourth-generation ABAC programming language developed by NextLabs to make implementation easier. Following the eXtensible Access Control Markup Language (XACML) standard, business users define declarative statements comprised of Subjects, Resources, Actions, and Conditions to capture local business concepts.



A policy set can mirror the language of NDAs, TIAs, PIAs, and other information sharing agreements and regulations. For example, ACPL supports exceptions to a parent policy rule to capture the structure of regulations, as well as supply highly structured, precise logic.



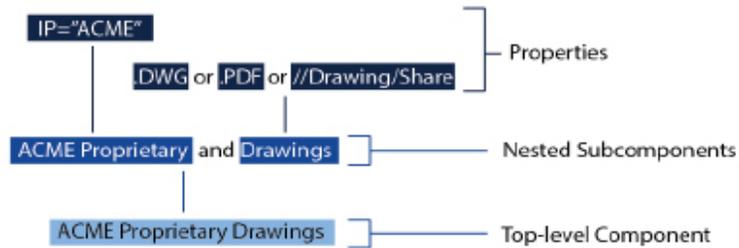
When required, policy can be written to target resource and subject relationships dynamically to capture business requirements that are defined more granularly. For instance, a single policy can elegantly protect all Product Team resources, no matter who is accessing them or where they are being accessed from.



Policy Component Model

Business users rely on the Policy Component Model to strategically tailor a component structure to reflect local business concepts and organizational hierarchies. The Policy Component Model also provides an intuitive “building block” approach to writing policies. Policy Components can be combined and nested using Boolean logic and then reused across multiple policies.

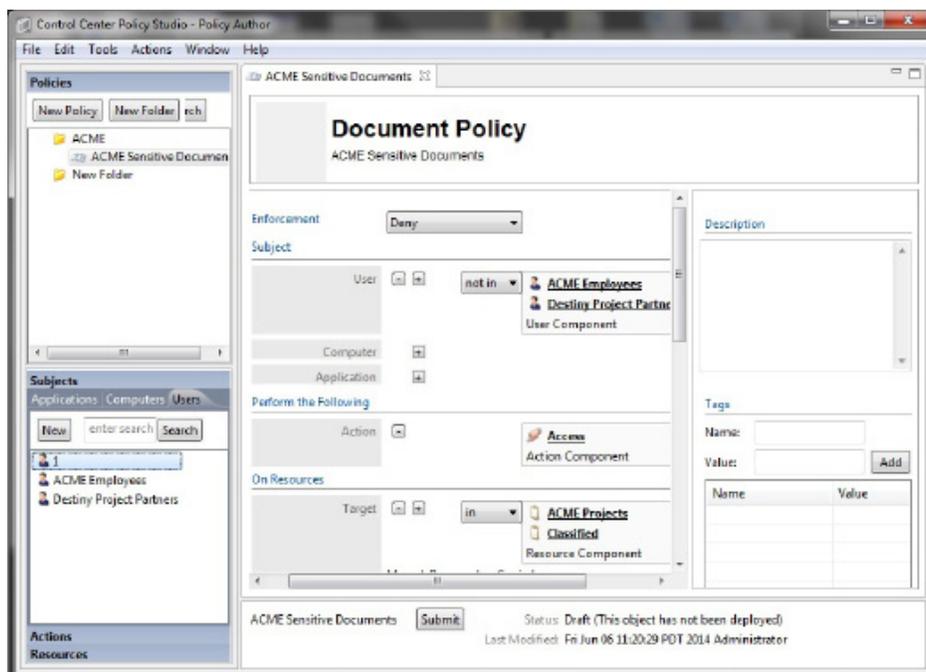
The Policy Component Model creates a layer of abstraction so policy designers—typically business users—need not be familiar with the technical details of underlying physical systems. Business concepts are expressed as user-defined properties of users and data (for example, citizenship, project team, clearance level, document classifications for circulation restriction or document status, file property, location, and so on).



Policy Lifecycle Management

Control Center’s graphic interface Policy Studio provides two policy development studios so Business Users—who are typically most familiar with information sharing requirements—can easily create and manage ACPL Policies.

Policy Author provides a drag-and-drop interface for creating policy. Business users simply drag and drop Policy Components into policy templates. Policy Manager applies a workflow to manage, deploy, and audit Policies and Policy Components, with automatic validation and segregation of duties. Within a single policy domain, policy administration can be delegated to different organizations, business units, departments, or teams.



Attribute Management

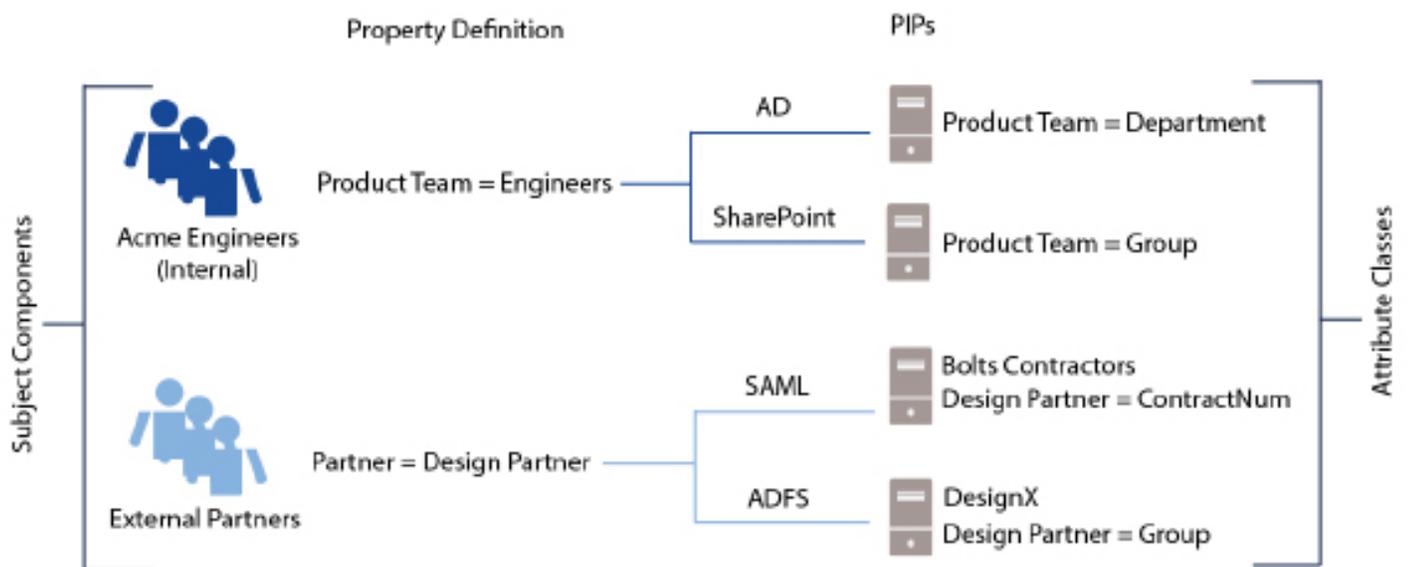
Attribute Management allows organizations to leverage attributes already available, both internally and externally, for consistent policy enforcement across the extended enterprise. Attributes available in disparate identity stores can be owned and administered by different teams (both internally and externally), at the same time as being centrally integrated with Control Center, so they are available during policy evaluation.

Attribute Management employs the following Control Center technologies (discussed in more detail below):

- **Attribute Mapping** – Define integrations with internal and external identity stores, without needing to “own” or be familiar with their business logic
- **Attribute Inventory** – Logically normalize attribute naming and storage conventions to reflect local business concepts
- **Policy Information Points (PIPs)** – Out-of-the-box support plus SDKs for the most common attribute stores in today’s extended enterprise

Attribute Mapping

Attribute mapping connects properties defined in the Policy Component model to attribute classes in attribute source systems (directories, databases, and applications). Attribute mapping does not require integrators be familiar with or “own” the actual attributes (which might be impossible in the extended enterprise). Instead, Attribute mapping establishes the location of each attribute store and the logic used to retrieve attributes.



Attribute Inventory

Today, IT organizations typically do not know which attributes are stored in an Identity and Access Management system, in different LOB applications, in groups, sites, and document libraries in Microsoft SharePoint Server, and other repositories. The Attribute Inventory applies structure to an organization's typically unstructured network information, so organizations know which users, groups, roles, data classification, and other attributes exist across their different systems.

The Attribute Inventory normalizes attributes from these disparate source systems and is aware of changes, such as when new users join the company when existing users changes roles, when new applications and data repositories need to be monitored when there are new integrations with external attributes stores, and so on.

Policy Information Points (PIPs)

Out-of-the-box, NextLabs supports integrations with common attribute stores which, once integrated, server as Policy Information Points (PIPs). Commonly used PIPs for identity attributes include LDAP and Active Directory, Microsoft SharePoint Server, HR applications, CRM applications, and Asset Management systems.

Business Policy Evaluation

In the extended enterprise, policy evaluation must be applied to data, no matter where it resides: on desktops, servers, and in the cloud. Control Center's Business Policy Evaluation provides data-centric evaluation, using the following core technologies (discussed in more detail below):

- **Policy Controller** – Cross-platform Policy Decision Point (PDP) that provides dynamic, attribute-based policy evaluation
- **Distributed Evaluation** – Online or offline, close to applications, desktops, and systems, or remotely for effortless scalability across the extended enterprise
- **Smart Deployment** – Performance optimization for dynamic enforcement with zero impact on user productivity.

Policy Controller

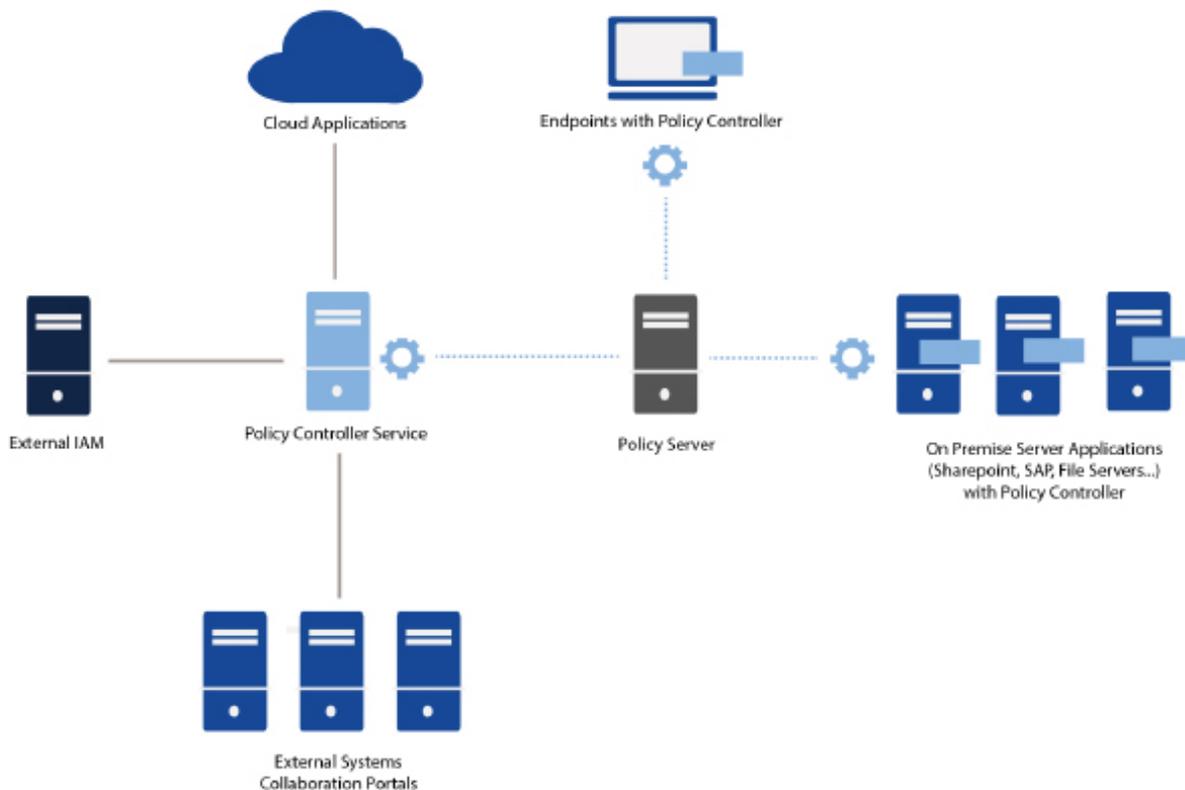
The Policy Controller is a PDP that runs anywhere to provide critical services, such as policy evaluation, configuration management, logging, security, tamper resistance, and obligation management. Upon an information access event, the Policy Controller can quickly evaluate relevant policies and return an appropriate response (allow, deny, or log).

Each Policy Controller uses digital certificates to authenticate with the Control Center service. All policy is encrypted and digitally signed by the Control Center and authenticated by the Policy Controller to ensure that only valid policies are activated.

Distributed Evaluation

Policy Controllers reside anywhere your data is to supply evaluation decisions to applications and systems where information access occurs. Policy Controllers can be co-located with server applications and systems to minimize network traversals, as well as installed on endpoint devices for policy enforcement that persists even when laptops are outside the network. Policy Controllers can run as a cloud service as well, returning decisions to authorization requests no matter where they originate across the extended enterprise.

As the following example implementation demonstrates, this architecture enables effortless scalability across the extended enterprise with high-availability inherently built into every evaluation and enforcement service.



Smart Deployment

Smart Deployment applies performance optimizations to policy evaluation, which removes any perceptible latency or interruptions to end user experience. Business Policies are optimized and packaged as policy bundles, which are custom-built for each Policy Controller and automatically contain the minimum set of policies relevant for each target host. Policy rules are also pre-evaluated so that only truly dynamic factors need to be evaluated in real time. Attribute information that is semi-static (that is, changes less than 5% a day) can be enrolled with Control Center, further reducing the amount of real-time attribute retrieval required.

Business Policy Enforcement

Wherever possible, automation should make information risk management easy for users: data owners, end users, auditors, and troubleshooters.

Business Policy Enforcement supplies automated services that replace the standard manual work of data centric security with the following Control Center technologies (each discussed in more detail below):

- **Policy Enforcers** – Automatically intercept information access requests and automate enforcement of decisions supplied by the Policy Controller.
- **Information Controls** – Allow data owners and business users to automatically apply controls to data, such as encryption, classification, access, and storage.
- **Centralized Audit** – Ubiquitous event logging gives auditors and troubleshooters visibility to access events across the extended enterprise.

Policy Enforcers

Policy Enforcers supply real-time interception of information access requests and serve as the Policy Enforcement Point (PEP) by routing information to Policy Controllers for evaluation. Control Center supplies several-of-the-box PEPs, as well as an SDK for Windows, Linux, and Unix environments.

Information Controls

Policy Enforcers provide automated information controls, which are flexible enough to adapt to levels of information risk. For example, at one end of the enforcement scale, monitoring events may be enough to mitigate risk. At the other end of the scale, an organization may need to block access, encrypt, or require a workflow.

A range of information controls automatically govern what users can do with data, including:

- **Classification** – Automatic or user-driven classification applies persistent metadata to files which can be used for policy evaluation, no matter where a file goes.
- **Encryption** – Automatically apply encryption to files based on classification, storage location, file attribute, or other data characteristic, or even the nature of the user action (for example, uploading to a public portal). Encryption persistently protects files even when they leave your domain.
- **User-Assisted Remediation** – Prompt users to complete tasks on files on their own, so data is properly secured before they share it. For example, users can be prompted to remove hidden data, verify email recipients, apply classifications to files.
- **Custom Workflow** – The extensible Obligation Manager framework can associate any custom executable with a policy enforcement event. Enforcement can be easily tailored to fit an organization's unique information control requirements.

Centralized Audit

Typically, system auditors and troubleshooters lack cross-system visibility to track data access events, resulting in many hours of manual work to troubleshoot or generate reports necessary to demonstrate compliance with information sharing policies and regulations. Control Center supplies ubiquitous event monitoring and logging for centralized audit across the extended enterprise.

Detailed attribute-based information can be collected and routed from all systems to the Control Center Activity Journal, a central collection point. Control Center Reporter, a web-based application that queries information stored in the Activity Journal, give administrators powerful access to information access and other system events. Reporter allows users to define, run, and share reports, and reports can be designed to track key data and user attributes relevant to local or industry-based sharing regulations and policies. Integration with third party reporting tools, system management, and



security event management systems is enabled through open web service interfaces.

Conclusion

ABAC allows for more granular access decisions. In order to implement data-centric controls that can handle the complex challenges of the extended enterprise, you need to classify your data, identify users, and create simple policies.

NextLabs Control Center is designed to supply these critical inputs and services to make the goal of data-centric ABAC possible – no matter where sensitive data is being accessed.

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

NextLabs, the NextLabs Logo, Compliant Enterprise, the Compliant Enterprise Logo, Deep Event Inspection, 360 Degree Enforcement, and ACPL are trademarks or registered trademarks of NextLabs, Inc. in the United States. All other trademarks are the property of their respective owners. 8-08.

© NEXTLABS INC. ALL RIGHTS RESERVED