

ISSA End-to-End Trust Working Group

ADAPTIVE ACCESS FRAMEWORK

CONCEPT, DEFINITION, APPLICATION

FEBRUARY 2012

Executive Summary

End-to-end trust for electronic transactions is vital to the continued healthy growth of an increasingly networked economy. Current trends in business computing have created potential for new classes of electronic business risk and transaction fraud. These new trends include:

1. Migration of business processes to external platforms including cloud & SAAS
2. Increased external collaboration: partners, consultants, external employees & customers
3. Business process support for consumer owned devices

The economic benefits for these new trends come from business process flexibility, infrastructure cost externalization, improved availability and improved user experience.

However, the impact of these trends in terms of increased business risk and compliance risk is straining the capabilities of existing end-to-end trust frameworks. Business owners require a more flexible business access control framework to balance risk and trust. The adaptive access framework presented in this paper provides an organic approach to adaptively balancing trust and risk. The adaptive access framework enables request by request evaluation of risk and the enforcement of combinations of claims to match the risk for that transaction. This enables trust assurance on an as-needed basis and provides a way to build end-to-end trust using attributes available within the infrastructure while minimizing impact to the end user experience.

Introduction

ISSA set up a cross-industry working group to gather and distribute problem statements, business cases, solutions guidance and operational experience around the core elements of end-to-end trust. End-to-end trust in this context is defined as the ability for the parties involved in an online transaction to have sufficient confidence in its security (Confidentiality, Integrity and Availability). There are three components of end-to-end trust when considered holistically help to establish trust in a transaction:

- Identity – of users, devices and services
- Conditions of trust – health and verification of the devices, network transmission paths, policy requirements (e.g., privacy, confidentiality, data integrity, data storage)
- Action – the type, the nature and impact of the transaction itself

This document aims to frame the end-to-end trust conversation in the context of a comprehensive framework for expressing, measuring and managing trust in online transactions. This document describes the need for a unified way to express various attributes of trust that help two parties involved in an online transaction to manage trust. It also aims to apply the overall framework to specific use cases involving a subset of attributes as a way to show how the framework helps solve business problems.

End-to-End Trust: Directions & Drivers

- **Business Risk:** All organizations are at risk of electronic fraud and data breach. Growth in external cloud-based business processes, consumer devices & mobile computing, and cross domain collaboration has opened the door for multiple new classes of business risk. Persistent attacks incorporating multiple attack vectors are straining the capabilities of existing fixed access control mechanisms.
- **Trusted Computing Technology:** There is growth in the availability of devices with hardware-based security capabilities. Hardware-based device identity, hardware measured device attributes and hardware-based remote attestation provided a path to increased end-to-end trust.
- **Increased operational flexibility:** Current trends in business have increased the range of operational scenarios which must be supported. Businesses require access control systems that can adapt to changes in usage patterns and adapt to changes in the threat landscape. An adaptive access framework can dynamically scale access requirements to balance transaction risk.
- **Business Benefits:** The Adaptive Access framework enables organizations to accomplish more with less:
 - **Increase end-to-end trust and flexible management of electronic business** risk through inclusion of a broader range of request attributes into access control decisions, including: device identity, device configuration and device health.
 - **Enhanced user experience** with adaptive frameworks which can transparently leverage device and transaction attributes to build trust and assurance without requiring additional actions from the end user.
 - **Manage cost and complexity** Adaptive access provides a way to augment existing access control systems and infrastructure so that trust assurance can be scaled to match risk. Increased trust can be achieved by incorporating device characteristics and transactions characteristic which already exist.

Market drivers for Adaptive Access

Traditional access management frameworks are based on fixed prescriptive mechanisms commonly centered on identify credentials, with centralized policy, centralized alerting and manual supervision. These fixed access control systems do not scale with the fluidity of today's business realities. Today's business fluidity is enabled through use of external business platforms, external collaboration and support for user owned devices. Each of these trends opens new categories of electronic business risk. Business owners require the ability to securely adapt these rapidly changing business trends and ever changing risk landscapes.

Adaptive access enables business managers to respond to electronic business access requests, based on a range of attributes related to the request. The attributes of a request are generally classified into four categories: Subject (Requestor), Context, Resource & Action. In contrast to systems which rely solely on identity, adaptive access increases end-to-end trust by including multiple factors to spread the decision risk. The framework is adaptive because an electronic business access decision can be based on the balance between attributes which quantify risk and attributes which quantify trust. The framework can adaptively require higher trust attributes in response to higher risk indicators. Identity is just one attribute among many which can be used to establish the balance between risk and trust.

Enhanced economics of business operations is the driver for adaptive access. Adaptive access does not impose worst case security requirements on all transaction requests. Instead risk is evaluated on a per transaction basis and demands for higher trust assurance are based on the risk of the specific transaction request. This dynamic costs/needs tradeoff translates into reduced need for security expenditure and reduced impact on the user experience.

The framework presented in this paper is a reference architecture which can be used as a guide to incorporate adaptive access into existing access frameworks. Initially, the focus is online transactions across multiple trust domains. The framework is based on policy components and a policy framework related to IETF RFC 3198; this enables standardized constructs that can be applied to non-standard implementations.

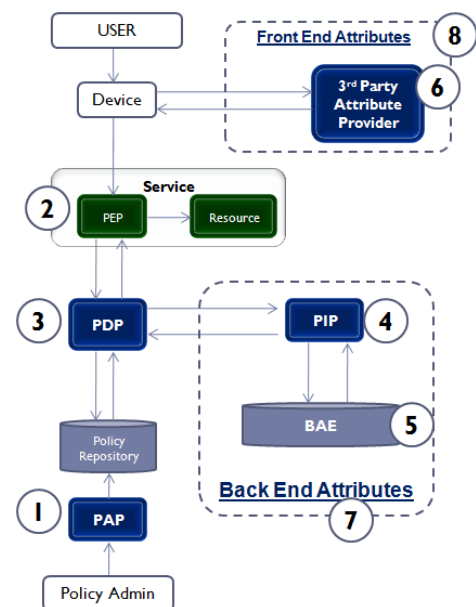
Adaptive Access Framework

The adaptive access framework is a reference architecture which leverages trusted attributes to enhance end-to-end trust in multiparty transactions without unnecessarily impacting the user experience or incurring excessive operational overheads. The framework can be used in a number of ways:

- A. To learn about trusted attributes
- B. To Learn how to incorporate trusted attributes into access control decisions
- C. To learn how to build adaptive access systems using trusted attributes

The framework is composed of the following functional blocks:

1. **Policy Administration Point (PAP):** A service that enables creation and management of policy
2. **Policy Enforcement Point (PEP):** The service responsible for making policy decision requests to the PDP.
3. **Policy Decision Point (PDP):** A service that is able to interpret the policy rules published by a PAP to make decisions to allow or deny requests.
4. **Policy Information Point (PIP):** A service which issues attributes requested by the PDP
5. **Back Attribute Exchange (BAE):** A service which provides Subject Attributes
6. **3rd Party Attribute Provider:** A service which provides attributes related to the subject
7. **Back End Attributes:** Attributes which are provided to the PDP via the PIP
8. **Front End Attributes:** Attributes which are provided to the PDP via the PEP



The standard sequence of events is (1) the policy enforcement point (PEP) sends all requests to the policy decision point (PDP) for processing. The framework represents the PEP & PDP as being disjoint components but they can be co-located in the same functional module, especially when low latency or high throughput is needed.

The function of the PDP is to make a decision regarding the request. (2) PDP makes its decision in accordance with policy. Part of the function of the PDP is to know how to retrieve applicable policy sets. Policy is organized in sets that are indexed using targets. Targets are simplified conditions for the subject, resources and action. The policy applies to a request when the policy target matches the request.

When the policy target matches the request, (3) the rules of the policy are evaluated, to determine how to disposition the request. The policy will specify the attributes of the request which are required to disposition the request. There are four classes of attributes: Resource, Subject, Action and Context. Conditions based on one or more attributes are at the heart of policy rules. If the condition is met (4) access is granted to the request.

Attribute Class	Sample Attribute
Subject	Name, Age, Sex, IsEmployee, EmployeeNumber, Account Number, PhoneNumber, SmartCard, IsManager
Resource	Resource Name, Service Level, Availability, Health
Action	Deposit, Query, Close, Read, Write
Condition	Request origination device identity, Origin IP_address, Originating Geo Location, Request time

The policy may also contain obligations which must be executed before the access is granted, frequently obligation are used to implement audit logging.

Attributes enable adaptive control

The framework proposed here is based on Attribute Based Access Control (ABAC). Access decisions are not based on the post-authentication rights of a subject. Access decisions are based on attributes of the request. The attribute-based access control policy specifies which claims need to be satisfied in order to grant access to an object. An example of a claim is "IsEmployee". Access is granted to users who can prove this claim. In the "IsEmployee" example the user may be anonymous as no proof of identity is required. Methods exist to support this form of access control.

The framework is based on ABAC but can be extended to Capability Based Access Control (CBAC) and Role Based Access Control (RBAC). Having attributes as the underlying decision variable introduces flexibility to adapt to changes in risk profile as can be seen in the following scenarios.

1. Scenario 1: Compromised Token Risk

Consider a service organization which has issued a hardware identity tokens to all members and an event arises where there is a possibility that the security of the tokens has been compromised.

- a. The potentially compromised security tokens introduce a new level of electronic business risk around identity of the subject.
- b. Adaptive access control provides flexibility for the organization to hedge against the token risk by adding other attributes which track with users.
- c. Machine identity of the requesting device might be a good example
 - i. There are a number of commercial solutions for accomplishing this, including Trusted Platform Module (TPM) based certificates.
- d. If it can be proven to the relying party that a transaction is coming from a device that is associated with the subject, then there is sufficient new trust which may cancel out the token risk.

2. Scenario 2: Insider Chargeback Fraud Risk

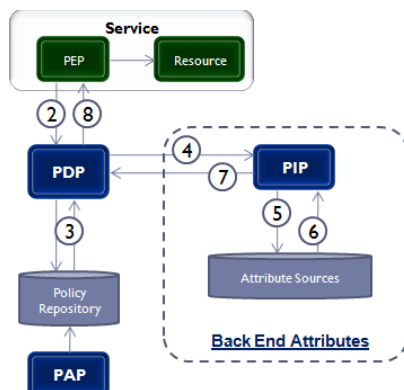
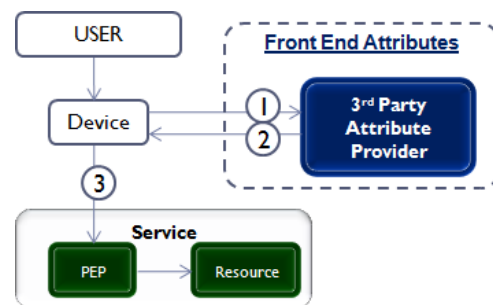
Consider an organization that is experiencing an excessive level of charge back fraud. To hedge against this risk the organizations restricts chargebacks over a certain limit to managers.

- a. This insider fraud scenario introduces risk around the trustworthiness of the subject.
- b. Adaptive access framework can hedge against this risk by adding an “iManager” attribute logically AND’ed with “Amount > Limit” term to the policy condition.
- c. Another policy with “Amount > Limit” in the condition will disposition chargeback requests for low risk amounts.

Internal Attributes Vs External Attributes

External attributes are referred to as “Front Channel Attributes” and these are attributes which are provided from a trust domain which is separate from the Relying Party trust domain. Front Channel attributes must be bound to the channel in such a way as to support the trust level required by the PEP of the relying party (RP).

Front Channel attributes are provided to the PDP along with the request, this may be accomplished in several ways including a hardware token recognized by the RP or as a security token signed by a Security Token Service (STS) which is trusted by the RP. There are several standard protocols which support this including: SAML, OpenId_connect, OAuth, Kerberos.



Internal attributes are referred to as ‘Back Channel Attributes’ and these are attributes which can be obtained from inside the trust domain of the relying party. The PDP requests back channel attributes through the Policy Information Point (PIP). Examples of back channel attributes include: Employee attributes stored in Active Directory or Requester attributes stored in the account database of a bank (i.e. Account Balance).

Commonly, access decisions are based on a combination of Front Channel & Back Channel attributes. Identity Federation is another

common scenario of claims based access where the access decision is federated back to the parent organization of the requesting party

Device Health and Device Configuration Attributes

The implication of cloud based business processes and consumer devices, is that customers, partners and employees potentially have access to potentially high-value resources anytime, from anywhere and from any device. This fluid access to corporate resources brings increased risk; potentially, users may request secure content from insecure devices. Adaptive access enables the business manager to create policies which identify high risk requests and includes device health attributes in policy rules.

Consider a medical supplies supply company that has to deal with individual patient requests and credit card transactions. When implementing a web service to automate ordering and billing for a large number of independent distributors and field service agents, multiple operational attributes may need to be considered to understand risk and trust before a request is dispositioned. The list can include but is not limited to:

1. Does the requested material contain Personally Identifiable Information (PII) or contain credit card information
2. What is known about the requesting device and can that data be trusted
3. Does the requesting device have up-to-date anti-virus and anti-malware
4. Does the requesting device have an encrypted hard drive
5. Does the geographic location from where this request originated have unique compliance requirements

In an adaptive access framework these questions can be formulated in terms of claims and the answers can be presented in terms of attributes. This enables real-time negotiation of trust. Access control policies can be designed to balance risk and trust on a request by request basis. The benefit of this adaptive control is that the whole system does not have to be designed for the worst case scenario. It may be that 60% of the requests are related to no-risk transactions, which require minimal security (i.e. downloads of marketing material).

This continuous adaptive balance between risk and trust has a number of benefits including:

1. Provides ease of use for these low-risk transactions with anywhere, anytime, any-device access.
2. Adaptively scales up security requirements when risky actions are requested.

When designing a system to support this scenario it is helpful to map out the required attributes and how they can be used. The following table illustrates attributes that may be used.

Attribute Class	Attribute	Attribute Source	Application
Target Resource	"SensitiveData" "Confidential"	Metadata tagged by: <ul style="list-style-type: none"> DLP Discovery Data Source Manual tagging 	To evaluate the risk associated with a request
Subject / Context	"DatVersion" "DatDate" "FW-Config"	NAC data repository using commercial 3rd IMV & IMC products.	Evaluate the risk associated with the target destination for down loaded data
Subject / Context	CorpManagedDev	Query AD via the PIP to check if this is a fully managed asset	Enable restrictions on "Confidential" data

Subject / Context	"FDE"	Acquired from the configuration management data base via PIP (e.g. PCCLM)	Ability to ensure continuous compliance for "SensitiveData"
Subject/Context	"StrongMachineID"	TPM-based PKI identity with Attestation	Ensure non-spoof-able & non-migrate-able requesting device identity for high-value information
Context	"IP-Address"	Packet transmission header	Verify whether this is an internal or external request.
Context	"TunnelType"	Packet Header	Access transmission risk
Context	"MsgAuthenticator"	Packet header	Access transport security

Why Adaptive Access Framework

The framework presented here provides an incremental path to implement higher levels of end-to-end trust. Changes in business usage and business architectures introduce a host of new security and risk concerns. Traditional enterprise domain centric approaches to access management are not flexible enough to keep up with market needs. The adaptive framework presented here enables a way to analyze risk on a per request basis and dynamically issue claims which match the specific request.

Conceptually the benefits of this adaptive access framework include:

1. **Control:** Provides a systematic organic mechanism to balance trust and risk on a per access request basis.
 - a. This avoids over provisioning of security and unnecessary user burden on low risk transactions.
2. **Security:** Adaptive access improves end-to-end security by incorporating a variety of independent factors into the access control decision.
 - a. User behavioral characteristics (i.e. Machine ID or Geo-location) can be incorporated into the decision
 - b. The health and configuration of the requesting device can also be used in the decision process if appropriate
3. **Incremental Adoption:** The framework outlines a progressive strategy for adding adaptive access capability to existing frameworks.
 - a. Existing access control investments can be leveraged and augmented

Strategically the adaptive access framework is aligned with industry initiatives to support business oriented architectures, including:

1. "Claims Based" access control which embraces and subsumes all of the access control strategies that have existed to-date.
2. Existing protocols including: SAML, OAuth and OpenId_connect
3. Externalization of security trends. The framework can be implemented as a gateway or a web agent which enables uniform access control to be extended to external web platforms.

From an implementation point of view, the framework can be configured using commercially available 3rd party products and 3rd party services.

1. There are a variety of PIP and PDP products currently available in the market
 - a. Axiomatics, Vordel, layer 7, Oracle

2. The framework is extensible to incorporate 3rd party attribute providers which currently exist in the market.
3. Endpoint management solutions, PCCLM solutions and MDM solutions are rich sources of device configuration and device health attributes, which can be integrated with the framework
4. Trusted computing technologies such as TPM and TNC can be used in the framework. Support for these Trusted technologies are available from a number of vendors including: Microsoft, Dell and Wave Systems

Planning for adaptive access

When a business depends on online exchange of value there are many important considerations, including:

1. Business risk and how that risk is likely to change over time
2. Compliance requirements
3. User experience
4. What available 3rd party products and 3rd party services that can be leveraged
5. Cost of building the access control system and cost of owning the access control system

The relative importance of these high-level considerations depends on the specifics to the business scenario. However there are some planning strategies which can be generally applied to many different scenarios, we will cover some of these general precepts here.

General design considerations

1. Legacy Tradeoffs: Adding adaptive capability to existing access control architecture can be constrained by the implementation of the existing system. In these scenarios it is best to ignore legacy-constraints when conceptualizing the high-level adaptive access design and record any subsequent adjustments and tradeoffs to accommodate legacy constraints.
2. External Users: Even when initial requirements are limited to internal users, always plan for external users that have no a priori relationship with the service provider.
3. Portability: Devote part of the planning process to portability; consider using a policy engine that is supported on many platforms. The other portability consideration is the ability to extend corporate access control to external platforms including cloud / SaaS. This may involve migration to a gateway or web agent implementation.
4. Fine Grained Access Control: Look for opportunities to externalize authorization wherever possible. Also leave room for a plan to harvest the business value that can be gained from audit logs of externalized authorization.
5. Compliance: Plan for reconciliation of provisioned privileges and actual privileges and plan for comparison of provisioned privileges and actual usage based on audit logs.
6. Upwards Delegation: Increasingly the business manager needs to be able to tune access control systems to changing business needs and changing threat landscapes. Incorporate graphical and/or high-level declarative management interfaces where possible.
7. Trusted Computing: Leverage hardware root-of-trust capabilities whenever possible, this includes TPM (Trusted Platform Module) for strong device identity and remote attestation of TPM measured attributes.

Going forward

This first version of the adaptive access framework provides guidance on how to incorporate adaptive access methods into existing access control systems, there is a lot more work which remains to be done. Future work to be done includes:

1. Inclusion of more access control scenarios
2. Guidelines for provisioning, issuance and lifecycle management of attributes
3. Framework for network of 3rd party attribute providers
4. Guidelines for modular policy and policy portability

If you are a professional working in the area of access control please come and join the ISSA End-to-End Working Group. To find out more or inquire about participation please contact the working group Co-Chair, Erik Bataller (embataller@gmail.com).