IBM Software
Thought Leadership White Paper

March 2012

Context-Enhanced Authorization

Improved security for data access via mobile devices

Contents

This white paper has been compiled by IBM Security Systems using research carried out together with Novay. It describes how Context-Enhanced Authorization can facilitate the increasingly sophisticated security policies needed to protect the mobile environment without hindering legitimate interaction.

1 Introduction

1.1 Maintaining security in an increasingly mobile environment

As mobile devices become increasingly common for performing both business and consumer activities, businesses have to find the right balance between protection from an ever-increasing range of security threats and capitalizing on the many advantages mobile devices offer.

Mobile devices such as laptops or mobile phones operate in potentially hostile environments with unsecured wireless connectivity and increased chance of theft or loss. The risk of fraudulent transactions using stolen, lost or hacked identities will considerably increase into the future. Simply controlling who can access what is rarely sufficient in such emerging environments. Additional measures are needed.

Context Enhanced Authorization (CEA) provides the necessary additional security without hampering the user experience. As the office becomes more a meeting venue than a place of work, and consumer interaction occurs on an increasingly ad-hoc basis, implementation of CEA opens up a safer world of possibilities.

The central idea behind Context Awareness is that aspects of the user's context such as time, place, method of access, recent behavior and relationships with others are taken into account while the user is interacting with online services. Applications can involve checking consistency with expected behavior (e.g. an advisor visiting a client is likely to access the organization's systems from the client's office) or enforcing the four-eyes principle (e.g. both users are inside the building, and in the same room).

Emerging risks associated to nomadic working scenarios for both employees and outsourcing can readily be mitigated by using CEA, which can be interpreted as an extension of the office security perimeter. CEA can help businesses meet strategic organizational goals (such as the ability to work from mobile phones) and maintain acceptable risk levels.

1.2 Comparison with Risk-Based Authentication

Until now, security processes have been more focused on Risk-Based Authentication (RBA), which works by recognizing factors such as a user's regular browser and computer. RBA is based upon access policies which are

pre-configured to static behavior and cannot be dynamically adapted to new constraints. This may stimulate the enforcement of all-or-nothing policies which prohibit users working away from the office from accessing certain critical applications. In an increasingly mobile world, these types of authorization policies may be too coarse and restrictive.

Context-Enhanced Authorization (CEA) facilitates much more fine-grained access policies based on context information which is determined by factors such as when and where users are and what they are up to. This information is automatically gathered from various mobile devices and wireless networks. With a greater matrix of factors in the authorization mix, CEA is more reliable for determining whether a risk level is acceptable. This results in authorization management that's both more flexible and more secure, making it possible to grant more access privileges without reducing the risk profile.

1.3 Context Parameters and their applications

Context-based security offers many advanced possibilities. Information gathered from a range of sources can be used to determine compliance within various context parameters.

- **Location.** Location information is either obtained directly from sensors like GPS or inferred from sensors like IP-address, Outlook Calendar, VPN gateway or GSM cell ID.
- **Social.** Social context can be used to determine who the user is with.
- **Activity.** Activity-based context is indicative of what a user is doing, e.g. working, travelling, or meeting.
- **Physiological.** Physiological context includes heart rate and body temperature and can be used to determine the user's mental state, i.e. is he/she scared (based on skin temperature measurements) or nervous (based on e.g. heart rate).

Monitoring context parameters over time can facilitate pattern or behavior recognition that can also be used as input for context-enhanced authorization.

The following table summarizes potential context sources and corresponding use cases, as applicable in a financial setting.

Context value	Context Source	Applicability	Risk mitigated
Location	IP-address, GPS, GSM cell ID, badge reader for physical access, or Outlook Calendar.	Is the bank employee at home, at the office or somewhere else? Allows for more fine-grained access control while not working in the office.	Prevent financial transactions while not in the office. Alert employee that confidentiality is at stake while working in a public place.
Togetherness (two employees)	Bluetooth and most of the above (fused) context sources	For 4-eyes cases or separation of duties.	Prevent fraud.
Proximity (between employee and device)	Inferred from location information sources.	Is the user in the proximity of the device conducting the transaction?	Prevent data loss in case the device is stolen.
Velocity	GPS	Is the employee working in the train?	Alert employee that confidentiality is at stake while working in the train, i.e. prevent data loss.
Emotion	Skin temperature sensor.	Is the employee threatened to conduct a transaction?	Risk of an intruder forcing an employee working at home to conduct a transaction.

Context value	Context Source	Applicability	Risk mitigated
Customer context	Customer call record.	Customer helpdesk employee has only access to customer information if the relevant customer has contacted him/her.	Prevent the customer helpdesk employee of unnecessary accessing arbitrary customer information.
Communication context	Type of device used, IP-address, or VPN connection context.	Is the employee working from a mobile device (that could very well be stolen) or from a remote location?	Prevent financial transactions conducted from a mobile device or from a remote location.
Behavior pattern	Inferred behavior from location, communication and other context sources (e.g. transaction log files)	Verify the actual identity of the user. This actually is risk-based authentication.	Prevent that someone else starts doing transactions on behalf of the employee.
Transaction context	The amount of money to be transferred, or only transactions to known or previously used bank accounts	Prevent large transactions from a mobile device or to unknown bank accounts.	Limit the set of possible financial transactions in the case that someone else has stolen the device of the employee.

Table 1: Overview of possible context sources that could be of relevance for context-enhanced authorization in a financial setting.

2 BENEFITS

2.1 Improved security

Context Enhanced Authorization uses both contextual and historical user information, along with data supplied during (Internet) transactions, to assess the probability of a user interaction being authentic or not. Aspects such as when (date and time), from where (location and IP address), how (what device), why and under what special conditions someone can access resources can be combined in the security mix. Historical user data may also be aggregated to determine user behavior and transaction patterns.

By constantly monitoring and analyzing context information, CEA makes it possible to maintain the desired security level and respond to new authorization constraints that may arise from changes in the situational context – e.g. a change in device, location or gap between activity. This ability to automatically adapt results in higher trustworthiness, security, usability, and flexibility.

The use of such contextual information can help optimize authorization functionality in a non-intrusive way. Based upon a risk profile and the amount of identity certainty, a rule engine can decide whether or not to grant access to services or allow transactions, in real time. If the risk is too high, additional step-up authentication can be asked or the authorization rights tightened.

2.2 Mobile and flexible

In addition to improving the level of security during transactions in need of authorization, context-enhanced authorization offers several other benefits:

- Adaptively limits access rights when using inherently untrusted platforms such as mobile phones
- Is user friendly, as it works behind the scenes.
- Can occur continuously during sessions and transactions that are specified as high risk (context may change during a session and may require new security settings).
- Is relatively cheap, as no expensive hardware is required and management overhead is minimal.

- Is easily extendable with new or other context parameters and as a consequence adaptive to new emerging threats/ risks.

3 CHALLENGES

Challenges in implementing CEA include authenticity of context, quality of context, privacy and complexity of context-aware policies.

3.1 Authenticity of Context

The success of Context Enhanced Authorization depends on overcoming a number of challenges such as false positives and accuracy of risk prediction. Context sources themselves are often not designed with security in mind, and can easily be misled to produce falsified context. Verification of the authenticity of the information provided is therefore essential, and this is achieved through several means.

- **Trusted sources.** The source of context information is the premier criterion for its credibility and quality. This can be a trusted subject/user, context provider or broker.
- **Cryptography.** Context-based digital signatures can be used to protect the authenticity (integrity) of the context information from trusted sources.
- **Distance bounding protocols.** Round-trip time measurements can be used to gain and verify location information about a user in different contexts. Drawbacks include that such location information is associated with a rather large area, there is considerable chance of error, and it works poorly in routed networks.
- **Proximity.** Being in the neighborhood of a trusted reference point is a common technique for proving user location. This is done by receipt of a security token known to the authenticator of the access controller which is transmitted from a nearby beacon, via location-limited short-range channels or limited-range radio broadcast.
- **Context history.** A user's context claim can be validated against a reference database containing history information of earlier claims made by users. If such a new claim is in line with what is expected based on the history database, it could be trusted. For example, the use of location history and movement patterns can improve location-tracking techniques. Furthermore, Kalman filtering can improve the accuracy of context information such as the current and future movements based on GPS-information. Despite the uncertainty of context information, probabilistic logic, fuzzy logic, Bayesian network and other mechanisms still allow reasoning about such information.
- **Comparison/uniqueness.** Possibly the most frequently used method for proving the authenticity of physical objects is simply comparing them to other authentic objects. As digital security and physical security become increasingly interdependent in context aware environments, this method can also be used to verify context information that is used for security purposes, e.g., context claims of several train travellers can be compared with each other. Logically, all travellers in the same train should have the same velocity. So if traveller X claims to be in train Y with velocity V1=60 km/hour and five other travellers also claim to be in train Y with velocity V2=75 km/hour, then it can be concluded that traveller X is faking his velocity and possibly his location.

3.2 Quality of context

Quality of Context (QoC) is the collective term for the certain amount of "vagueness" of context.

- **Incorrect information.** Sources (i.e., sensors) providing context information may be incorrect - whether by accident or on purpose due to an active attacker - with a certain probability.
- **Precision.** Sources will have a certain precision (accuracy, granularity) with which they describe the real world situation.
- **Delay.** There can be time gaps between gathering and delivery of context information.
- **Temporal resolution.** Context information becomes less relevant beyond a certain time frame.
- **Spatial resolution.** There may be limits to the physical area within which context information is applicable.
- **Separation between device and user.** If the context source is based on a mobile device which the user is always supposed to be carrying, the mobile device may not be as tightly coupled to the user as desired (i.e., the user may misplace the device).

Accuracy is typically improved by fusing context sources with different QoC results. When inferring, context based on measured data probabilities should be taken into account. In context-aware authorization rules the QoC of the context should be factored in. This can add to the complexity of the rules, e.g., allow a certain transaction if the employee has >95% probability of being <50 meters from their house. But what do you conclude if an employee has 99% probability

of being <100 meters from their house? QoC also covers privacy sensitiveness of context information, because the more privacy sensitive the information is, the better quality it is likely to be.

3.3 Privacy

The use of context information may come at the cost of some of the user's privacy, with centralized components (e.g., the context sensing infrastructure, the authorization policy engine) having the opportunity to build up an enormous amount of privacy-sensitive context data on individual subjects.

In general, privacy issues can be dealt with by giving users control over what data is collected and making the purposes for collection transparent. However, depending on the type of organization, such user control may not be workable. Local rules and regulations may not permit collection of employee context data for "security purposes". It may be compulsory to give individual employees a choice between participating in (and having privacy sensitive data being collected) or objecting to the system.

The quality indicators (see 4.2 QUALITY OF CONTEXT) for context information can also be used to indicate how privacy-sensitive different types of information are. For example, the information that Alice is in Amsterdam is less privacy-sensitive than the information that she is at Central Station. Both pieces of context present her location, but the latter is more precise. The Principle of Least Privilege states that a subject should be allowed the least possible access that is required to fulfill their responsibilities. Applying this principle to users' context information requires empowering the user to have control over the quality of context information that is collected, stored and communicated.

3.4 Complexity of Context-Aware Policies

When context parameters are added to access policy rules, these rules tend to become more complex. It may become harder to ensure that the produced policies are complete, safe, and conflict-free. It may also get harder to determine which rules are no longer necessary.

Whether this is an actual problem depends on the abstraction level at which contextual aspects of policies can be expressed. If this abstraction level is sufficiently high, the policy rules may be self-explanatory. On the other hand, if the abstraction level is too high it becomes harder to reason about the authenticity and quality of the context sources producing the context information. Authors of authorization policies have to take these aspects into account as they influence the risk of unauthorized access.
In order to make the access control process truly dynamic and transparent, context aware security policies need to be both handcrafted in advance and generated on the fly. Introducing context parameters as part of the policy ('user X has access to service Y if he is in the building') and using the access control model learned over time to functionally adapt those policies should result in more robust, flexible and scalable access control solutions.

4 ARCHITECTURE

4.1 ABAC and XACML

Context information can be considered as attributes that are communicated during an authorization session. The attribute based access control (ABAC) paradigm facilitates this.
In ABAC, attributes are used to convey authorization information from (central) authorities to relying parties.
An important (perhaps, the de-facto) standard describing the architecture, syntax, and protocols for ABAC is XACML.

The data flow model of XACML follows the pull model of the authorization decision query. Briefly, the access request is received by the policy enforcement point (PEP) which then communicates it to the decision point (PDP). The PDP evaluates the access request with regards to the applicable policy set, policy or rule and replies with an authorization decision. In order to make a decision, the PDP obtains attributes associated with the client issuing the access request, the resource that is being accessed and the environment in which the access request is taking place. Such attributes may very well be contextual attributes and are retrieved from the policy information points (PIPs). The pull model of XACML makes it less suitable for dealing with dynamically changing context parameters such as location during a user session. Enforcing access on a transaction basis, however, can be pull-based and therefore be carried out via XACML.
The XACML architecture is shown in Figure 1.

In XACML, requests, responses, and policies may contain references to so-called environment attributes that can be used to specify context. Built-in attributes are current date

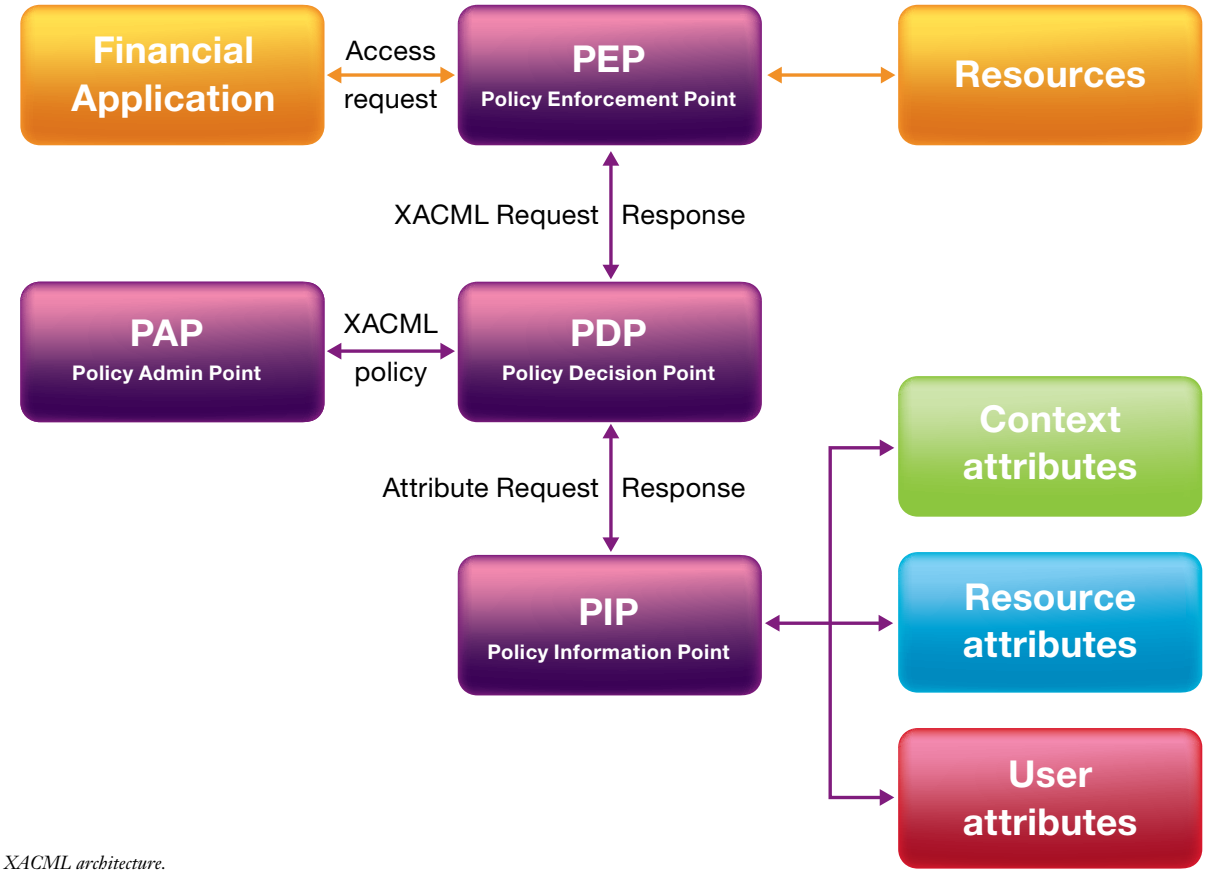

Figure 1: XACML architecture.

and time, but other attributes can be defined as well. ABAC, which due to the move to the cloud is at the center of attention, offers an important opportunity to start implementing context-enhanced authorization. The XACML policy framework standard is best equipped to implement an ABAC model based on context attributes.

Context-enhanced authorization is implemented by connecting context sensors to an ABAC based identity and access management (IAM) system, which functions as a PDP (see Figure 1). Typically the sensors are part of a context measuring infrastructure that functions as a PIP to the IAM system. This infrastructure is responsible for the discovery, collection, interpretation and communication of context information. To fully exploit the possibilities that context offers, the infrastructure can combine ("fuse") context information of different types, from different sensors (or the same sensor, but collected at different levels of abstraction) and enrich the context information by applying statistical analysis o n the measured context based on historical measurements.

Different architectural choices can be made:

- The abstraction level of inferred context information (i.e., how "smart" should the infrastructure be?).
- The trigger event model for collecting and communicating context information (i.e., pull, based on authorization events, or, push, based on context events).
- The openness of the context infrastructure (i.e. what are the context sources that are used and how trustworthy are they?).

These architectural choices have significant consequences for the (non-functional) requirements that should be put both on the infrastructure and on the IAM system.

5. About IBM's Security Portfolio

5.1 IBM

IBM's security portfolio provides security intelligence that helps organizations holistically protect their people, infrastructure, data and applications. IBM offers solutions for identity and access management, database security, application development, risk management, endpoint management, network security and more. IBM operates the world's broadest security research, development and delivery organization. This comprises nine security operations centers, nine IBM Research centers, eleven software security development labs and an Institute for Advanced Security with chapters in the United States, Europe and Asia Pacific. IBM monitors 13 billion security events per day in more than 130 countries and holds more than 3,000 security patents.

5.2 Tivoli Security Policy Manager

To facilitate Context-Enhanced Authorization as described in this document, IBM offers Tivoli® Security Policy Manager (TSPM), which helps implement context-aware authentication by centralizing security policy management and providing fine-grained data access control for applications, databases, portals and services.

5.3 For more information

A demo of context enhanced authorization using TSPM is currently available upon request. Please contact your IBM account manager for more information or visit: **ibm.com**/software/nl/tivoli

IBM Nederland B.V.
Johan Huizingalaan 765
1066 VH Amsterdam
The Netherlands

Produced in The Netherlands
March 2012