COMPUTING SCIENCE

A Novel Approach to Access Control for the Web

Machulak, M., van Moorsel, A.

TECHNICAL REPORT SERIES

No. CS-TR-1157 July, 2009

TECHNICAL REPORT SERIES

No. CS-TR-1157 July, 2009

A Novel Approach to Access Control for the Web

M Machulak, A van Moorsel

Abstract

The rapidly developing Web environment provides users with a wide set of rich services as varied and complex as desktop applications. Those services are collectively referred to as "Web 2.0", with examples such as Google Docs, Wikipedia, Wordpress or Flickr, that allow users to create, manage and share their content online. By switching from desktop applications to their Web equivalents more and more data gets released online. It is the user who creates data, who shares and disseminates this data, and who accesses it. Storing and sharing resources over a highly collaborative "Web 2.0" environment poses new security challenges. Access control, in particular, is currently poorly addressed in such an environment and is not well suited to the increasing amount of resources that are available online. We propose a novel approach to access control for the Web. Our approach puts a user in full control of their resources which may be scattered across multiple Web applications. Unlike existing authorization systems, it relies on a user's centrally located security requirements for those resources.

Printed and published by the University of Newcastle upon Tyne,
Computing Science, Claremont Tower, Claremont Road,
Newcastle upon Tyne, NE1 7RU, England.

Bibliographical details

MACHULAK, M., VAN MOORSEL, A.

A Novel Approach to Access Control for the Web
[By] M Machulak, A van Moorsel

Newcastle upon Tyne: University of Newcastle upon Tyne: Computing Science, 2009.

(University of Newcastle upon Tyne, Computing Science, Technical Report Series, No. CS-TR-1157)

Added entries

UNIVERSITY OF NEWCASTLE UPON TYNE
Computing Science. Technical Report Series. CS-TR-1157

Abstract

The rapidly developing Web environment provides users with a wide set of rich services as varied and complex as desktop applications. Those services are collectively referred to as “Web 2.0”, with examples such as Google Docs, Wikipedia, Wordpress or Flickr, that allow users to create, manage and share their content online. By switching from desktop applications to their Web equivalents more and more data gets released online. It is the user who creates data, who shares and disseminates this data, and who accesses it. Storing and sharing resources over a highly collaborative “Web 2.0” environment poses new security challenges. Access control, in particular, is currently poorly addressed in such an environment and is not well suited to the increasing amount of resources that are available online. We propose a novel approach to access control for the Web. Our approach puts a user in full control of their resources which may be scattered across multiple Web applications. Unlike existing authorization systems, it relies on a user’s centrally located security requirements for those resources.

About the author

Aad van Moorsel joined the University of Newcastle in 2004. He worked in industry from 1996 until 2003, first as a researcher at Bell Labs/Lucent Technologies in Murray Hill and then as a research manager at Hewlett-Packard Labs in Palo Alto, both in the United States.

Aad got his PhD in computer science from Universiteit Twente in The Netherlands (1993) and has a Masters in mathematics from Universiteit Leiden, also in The Netherlands. After finishing his PhD he was a postdoc at the University of Illinois at Urbana-Champaign, Illinois, USA, for two years.

Aad has worked in a variety of areas, from performance modelling to systems management, web services and grid computing. In his last position in industry, he was responsible for HP's research in web and grid services, and worked on the software strategy of the company.

His research agenda aims at establishing an intelligent enterprise. The goal is to establish objective, quantitative methods to improve IT-related decision making, eventually fully automated. This involves mathematical modelling and algorithms as well as service-oriented software implementations. DTI, EPSRC and EU-funded collaborations are ongoing with Hewlett-Packard, Merrill-Lynch, Warwick, Bath, UCL, various universities throughout Europe, and the Business School in Newcastle.

Maciej Machulak received his MSc in Computing Engineering from Wroclaw University of Technology in Poland in 2007. During his studies he was an Erasmus Exchange Student at Newcastle University. Maciej Machulak additionally completed the Advanced MSc degree in "System Design for Internet Applications" (SDIA) in 2007 at Newcastle University and his thesis was awarded Best 2007 SDIA Thesis. This degree included an industrial placement at Red Hat UK Ltd. Maciej's main task was to develop a framework for transactional Web Services. Before commencing his PhD studies Maciej was also employed as an intern at Red Hat and worked on embedded tools for transaction monitoring inside JBoss Application Server.

Maciej Machulak is currently a PhD student working with Dr. Aad van Moorsel on the Trust Economics project, funded by the UK Technology Strategy Board (TSB). Maciej's project is concerned with building a novel authorization system for Web environment which allows users to flexibly adapt access control for their Web resources to their particular security requirements.

Suggested keywords
ACCESS CONTROL, WEB

A Novel Approach to Access Control for the Web

Maciej Machulak
PhD Student
School of Computing Science
University of Newcastle upon Tyne
Newcastle upon Tyne,
NE1 7RU, UK
m.p.machulak@ncl.ac.uk

Aad van Moorsel
Reader
School of Computing Science
University of Newcastle upon Tyne
Newcastle upon Tyne,
NE1 7RU, UK
aad.vanmoorsel@ncl.ac.uk

I. Introduction

The rapidly developing Web environment provides users with a wide set of rich services as varied and complex as desktop applications. Those services are collectively referred to as "Web 2.0", with examples such as Google Docs, Wikipedia, Wordpress or Flickr, that allow users to create, manage and share their content online. By switching from desktop applications to their Web equivalents more and more data gets released online. It is the user who creates data, who shares and disseminates this data, and who accesses it.

Storing and sharing resources over a highly collaborative "Web 2.0" environment poses new security challenges. Access control, in particular, is currently poorly addressed in such an environment and is not well suited to the increasing amount of resources that are available online. We propose a novel approach to access control for the Web. Our approach puts a user in full control of their resources which may be scattered across multiple Web applications. Unlike existing authorization systems, it relies on a user's centrally located security requirements for those resources.

II. Problem statement

In contrast to desktop systems, the "Web 2.0" environment does not allow use of a single access control (authorization) mechanism, single policy language or single management tool for various Web applications. Authorization is often tightly bound to a Web application and has limited flexibility in terms of its configuration or adaptation to particular user's security requirements [1]. Access control policies are composed in incompatible policy languages and are maintained separately at every Web application. Heterogeneity and distribution of those policies poses problems in introducing new rules or modifying existing ones. More advanced custom access control solutions require modifications on the client side and cannot be easily adopted on a larger scale.

Recently proposed approaches do not fully address shortcomings of authorization in the "Web 2.0" environment. Lockr [2] is an access control system based on social relationships that separates content delivery and sharing from managing social networking information. Users are able to maintain a single social network and base Access Control Lists for different applications on social relationships. The Menagerie system [3] facilitates organization and sharing of collections of Web service objects. Such heterogeneous Web services expose access operations through a well-defined API and can be mounted into a local file system namespace.

Neither of the described approaches addresses the previously identified shortcomings in authorization for the Web. Lockr focuses on social relationships only. It does not define a generic model of authorization that can be plugged in to various Web applications and provide arbitrarily complex access control for heterogeneous and distributed resources. Menagerie, on the other hand, restricts access to Web resources through its Menagerie File System component.

III. Approach

We propose a novel user-centric authorization system for the "Web 2.0" environment that is able to address all of the previously identified shortcomings in existing authorization solutions. Our proposal puts a user in full control of their resources and relies on a user's centrally located security requirements for those resources.

A. Proposed Architecture

Our system is built on the concept of centrally expressed user's security requirements that are applied to a user's distributed Web resources. Such security requirements are expressed in form of access control policies and are stored and evaluated in a specialized component. We refer to this component as Security Provider (SP). A user delegates access control functionality from their Web applications to a Security Provider using a well-defined API. Architecture of the proposed system is depicted in Fig. 1.

A Security Provider allows a user to define access control policies for their online resources in a uniform way irrespective of the Web application that hosts those resources. It makes access control decisions based on those policies. It provides

Submitted for poster presentation at IEEE Symposium on Security and Privacy 15 April 2009. Accepted 20 April 2009. Final version submitted 15 May 2009. Poster presented 18 May 2009.

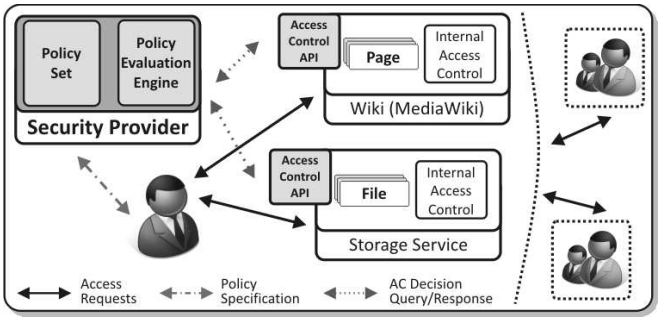

Fig. 1. User-centric access control system for the Web.

functionality of policy administration and policy decision points, such as those specified in [4].

In our proposal a particular Security Provider is chosen and controlled by a user. We base our concept on that used in OpenID [5] where a user chooses their preferred Identity Provider according to their requirements and preferences. In our system such requirements may include available policy languages, policy editors or policy management tools.

Access control functionality of a Security Provider is used by a user's chosen set of Web applications. Those applications delegate authorization using a well-defined API and are only concerned with enforcing access control decisions.

Establishing trust relationships between Web applications and a Security Provider and composing security policies is done with the use of a Web browser. Interactions between components of the architecture use HTTP as the transport protocol and conform to the REST architectural style.

B. Interactions

In our system a user interacts with the Web as usual and must only complete two simple steps of an initialization phase. At first, a user registers with their preferred Security Provider. Secondly, a user configures their chosen set of Web applications to delegate access control for all or part of a user's resources to this component. Such configuration involves providing a URL of a user's Security Provider.

A user interacts with their Web applications in a usual way by creating new data or uploading existing data. When access control needs to be applied to a resource, a user is redirected to their preferred SP. Either new rules are specified or an already composed policy is applied to a resource.

Access requests to Web resources are performed as usual. Authorization decisions are made by a Security Provider and are enforced by Web applications. Communication between Web applications and a Security provider is based on a simple request/response protocol and is transparent to a user.

C. Advantages

We recognize numerous advantages of our user-centric access control system for the Web:

1) A user is empowered to express their security requirements in a uniform way and store them in form of access control policies in a central location. Those policies can be applied to heterogeneous and distributed resources hosted by different Web applications;
2) Access control rules are composed and stored centrally which gives a user means to easily introduce new rules or change existing ones with minimum effort. Additionally, a user is given a consolidated view of the applied security mechanisms across multiple Web applications;
3) Management of access control policies is simplified through policy reusability. When a resource is moved from one Web application to another then the same policy can be applied with minimum effort;
4) A user can choose their preferred Security Provider which meets a user's requirements in terms of available policy languages, policy editors or policy management tools;
5) Our architecture does not require any modifications on the client side and relies on well-accepted technologies.

IV. Contributions

Our work presents a novel authorization system that fits precisely in the highly collaborative "Web 2.0" environment. In our system a user creates and shares content as before using various Web applications, but is additionally empowered to flexibly control access to their increasing amount of online resources. Permissions to data are managed in a central location and applied to resources which can be scattered across multiple Web applications.

We externalize authorization from Web applications and encapsulate it in form of an easily pluggable Web service. With this we aim to show that not only functional but also non-functional parts of Web applications should be provided in the form of services that work in a similar way to the Web itself. We believe that there is scope for providers of such services. Well-defined interaction sequences and interfaces of such services as the proposed access control should facilitate their adoption among Web applications.

V. Progress and Future Work

We have implemented a prototype Security Provider. We are currently implementing authorization plugins for the Wordpress blogging platform, the MediaWiki software that is used by Wikipedia, and our custom Web storage service.

Our immediate plans include defining the API of the SP component and the access control API for Web applications. Additionally, we plan to define the interaction sequences between different components of the proposed architecture.

References

[1] M. Hart, R. Johnson, and A. Stent, "More content - less control: Access control in the web 2.0," in *WOSP '08: Proceedings of the first workshop on Online social networks*. New York, NY, USA: ACM, 2008, pp. 43–48.

[2] A. Tootoonchian, K. K. Gollu, S. Saroiu, Y. Ganjali, and A. Wolman, "Lockr: social access control for web 2.0," in *WOSP '08: Proceedings of the first workshop on Online social networks*. New York, NY, USA: ACM, 2008, pp. 43–48.

[3] R. Geambasu, C. Cheung, A. Moshchuk, S. D. Gribble, and H. M. Levy, "Organizing and sharing distributed personal web-service data," in *WWW '08: Proceeding of the 17th international conference on World Wide Web*. New York, NY, USA: ACM, 2008, pp. 755–764.

[4] “OASIS eXtensible Access Control Markup Language (XACML),” http://www.oasis-open.org./committees/xacml/, 2005, version 2.0.
[5] “OpenID Specifications,” http://openid.net/developers/specs/, 2007.