

Network Working Group
Request for Comments: 3837
Category: Informational

A. Barbir
Nortel Networks
O. Batuner
Independent consultant
B. Srinivas
Nokia
M. Hofmann
Bell Labs/Lucent Technologies
H. Orman
Purple Streak Development
August 2004

Security Threats and Risks for Open Pluggable Edge Services (OPES)

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

The document investigates the security threats associated with the Open Pluggable Edge Services (OPES) and discusses the effects of security threats on the underlying architecture. The main goal of this document is threat discovery and analysis. The document does not specify or recommend any solutions.

Table of Contents

- 1. Introduction 2
- 2. OPES Data Flow Threats 4
 - 2.1. OPES Flow Network Level Threats 5
 - 2.1.1. Connection-Flow Denial-of-Service (DoS). 6
 - 2.1.2. Threats to Network Robustness. 6
 - 2.2. OPES Flow Application Level Threats. 6
 - 2.2.1. Unauthorized OPES Entities 6
 - 2.2.2. Unauthorized Actions of legitimate OPES Entities 7
 - 2.2.3. Unwanted Content Transformations 7
 - 2.2.4. Corrupted Content 7
 - 2.2.5. Threats to Message Structure Integrity 8
 - 2.2.6. Granularity of Protection 8
 - 2.2.7. Risks of Hop-by-Hop Protection 8
 - 2.2.8. Threats to Integrity of Complex Data 8
 - 2.2.9. Denial of Service (DoS) 9
 - 2.2.10. Tracing and Notification Information 9
 - 2.2.11. Unauthenticated Communication in OPES Flow 9
- 3. Threats to Out-of-Band Data 9
 - 3.1. Threats that Endanger the OPES Data Flow 10
 - 3.2. Inaccurate Accounting Information 10
 - 3.3. OPES Service Request Repudiation 11
 - 3.4. Inconsistent Privacy Policy 11
 - 3.5. Exposure of Privacy Preferences 11
 - 3.6. Exposure of Security Settings 11
 - 3.7. Improper Enforcement of Privacy and Security Policy . . 11
 - 3.8. DoS Attacks 12
- 4. Security Considerations 12
- 5. References 12
 - 5.1. Normative References 12
 - 5.2. Informative References 12
- 6. Acknowledgements 12
- 7. Authors' Addresses 13
- 8. Full Copyright Statement 14

1. Introduction

The Open Pluggable Edge Services (OPES) [1] architecture enables cooperative application services (OPES services) between a data provider, a data consumer, and zero or more OPES processors. The application services under consideration analyze and possibly transform application-level messages exchanged between the data provider and the data consumer. The OPES processor can distribute the responsibility of service execution by communicating and collaborating with one or more remote callout servers. The details of the OPES architecture can be found in [1].

Security threats with respect to OPES can be viewed from different angles. There are security risks that affect content consumer applications, and those that affect the data provider applications. These threats affect the quality and integrity of data that the applications either produce or consume. On the other hand, the security risks can also be categorized into trust within the system (i.e., OPES service providers) and protection of the system from threats imposed by outsiders such as hackers and attackers. Insiders are those parties that are part of the OPES system. Outsiders are those entities that are not participating in the OPES system.

It must be noted that not everyone in an OPES delivery path is equally trusted. Each OPES administrative trust domain must protect itself from all outsiders. Furthermore, it may have a limited trust relationship with another OPES administrative domain for certain purposes.

OPES service providers must use authentication as the basis for building trust relationships between administrative domains. Insiders can intentionally or unintentionally inflict harm and damage on the data consumer and data provider applications. This can be through bad system configuration, execution of bad software or, if their networks are compromised, by inside or outside hackers.

Depending on the deployment scenario, the trust within the OPES system is based on a set of transitive trust relationships between the data provider application, the OPES entities, and the data consumer application. Threats to OPES entities can be at the OPES flow level and/or at the network level.

In considering threats to the OPES system, the document will follow a threat analysis model that identifies the threats from the perspective of how they will affect the data consumer and the data provider applications.

The main goal of this document is threat discovery and analysis. The document does not specify or recommend any solutions.

It is important to mention that the OPES architecture has many similarities with other so called overlay networks, specifically web caches and content delivery networks (CDN) (see [2], [4]). This document focuses on threats that are introduced by the existence of the OPES processor and callout servers. Security threats specific to content services that do not use the OPES architecture are considered out-of-scope of this document. However, this document can be used as input when considering security implications for web caches and CDNs.

The document is organized as follows: Section 2 discusses threats to OPES data flow on the network and application level, section 3 discusses threats to other parts of the system, and section 4 discusses security considerations.

2. OPES Data Flow Threats

Threats to the OPES data flow can affect the data consumer and data provider applications. At the OPES flow level, threats can occur at Policy Enforcement Points, and Policy Decision Points [3], and along the OPES flow path where network elements are used to process the data.

A serious problem is posed by the very fact that the OPES architecture is based on widely adopted protocols (HTTP is used as an example). The architecture document specifically requires that "the presence of an OPES processor in the data request/response flow SHALL NOT interfere with the operations of non-OPES aware clients and servers". This greatly facilitates OPES' deployment, but on the other hand a vast majority of clients (browsers) will not be able to exploit any safeguards added as base protocol extensions.

For the usual data consumer, who might have questions such as (Where does this content come from? Can I get it another way? What is the difference? Is it legitimate?). Even if there are facilities and technical expertise present to pursue these questions, such thorough examination of each result is prohibitively expensive in terms of time and effort. OPES-aware content providers may try to protect themselves by adding verification scripts and special page structures. OPES-aware end users may use special tools. In all other cases (non-OPES aware clients and servers) protection will rely on monitoring services and investigation of occasionally discovered anomalies.

An OPES system poses a special danger as a possible base for classical man-in-the-middle attacks. One of the reasons why such attacks are relatively rare is the difficulty in finding an appropriate base: a combination of a traffic interception point controlling a large flow of data and an application codebase running on a high-performance hardware with sufficient performance to analyze and possibly modify all passing data. An OPES processor meets this definition. This calls for special attention to protection measures at all levels of the system.

Any compromise of an OPES processor or remote callout server can have a ripple effect on the integrity of the affected OPES services across all service providers that use the service. To mitigate this threat, appropriate security procedures and tools (e.g., a firewall) should be applied.

Specific threats can exist at the network level and at the OPES data flow level.

2.1. OPES Flow Network Level Threats

OPES processor and callout servers are susceptible to network level attacks from outsiders or from the networks of other OPES service providers (i.e., if the network of a contracted OPES service is compromised).

The OPES architecture is based on common application protocols that do not provide strong guarantees of privacy, authentication, or integrity. The IAB considerations [4] require that the IP address of an OPES processor be accessible to data consumer applications at the IP addressing level. This requirement limits the ability of service providers to position the OPES processor behind firewalls and may expose the OPES processor and remote callout servers to network level attacks. For example, the use of TCP/IP as a network level protocol makes OPES processors subject to many known attacks, such as IP spoofing and session stealing.

The OPES system is also susceptible to a number of security threats that are commonly associated with network infrastructure. These threats include snooping, denial of service, sabotage, vandalism, industrial espionage, and theft of service.

There are best practice solutions to mitigate network level threats. It is recommended that the security of the OPES entities at the network level be enhanced using known techniques and methods that minimize the risks of IP spoofing, snooping, denial of service, and session stealing.

At the OPES Flow level, connection-level security between the OPES processor and callout servers is an important consideration. For example, it is possible to spoof the OPES processor or the remote callout server. There are threats to data confidentiality between the OPES processor and the remote callout server in an OPES flow.

The next subsections cover possible DoS attacks on an OPES processor, remote callout server or data consumer application, and network robustness.

2.1.1. Connection-Flow Denial-of-Service (DoS)

OPES processors, callout servers, and data consumer applications can be vulnerable to DoS attacks. DoS attacks can be of various types. One example of a DoS attack is the overloading of OPES processors or callout servers by spurious service requests issued by a malicious node, which denies the legal data traffic the necessary resources to render service. The resources include CPU cycles, memory, network interfaces, etc. A Denial-of-Service attack can be selective, generic, or random in terms of which communication streams are affected.

Distributed DoS is also possible when an attacker successfully directs multiple nodes over the network to initiate spurious service requests to an OPES processor (or callout server) simultaneously.

2.1.2. Threats to Network Robustness

If OPES implementation violates end-to-end addressing principles, it could endanger the Internet infrastructure by complicating routing and connection management. If it does not use flow-control principles for managing connections, or if it interferes with end-to-end flow control of connections that it did not originate, then it could cause Internet congestion.

An implementation that violates the IAB requirement of explicit IP level addressing (for example, by adding OPES functional capabilities to an interception proxy) may defeat some of the protective mechanisms and safeguards built into the OPES architecture.

2.2. OPES Flow Application Level Threats

At the content level, threats to the OPES system can come from outsiders or insiders. The threat from outsiders is frequently intentional. Threats from insiders can be intentional or accidental. Accidents may result from programming or configuration errors that result in bad system behavior.

Application level problems and threats to the OPES systems are discussed below:

2.2.1. Unauthorized OPES Entities

Although one party authorization is mandated by the OPES architecture, such authorization occurs out-of-band. Discovering the presence of an OPES entity and verifying authorization requires special actions and may present a problem.

Adding notification and authorization information to the data messages (by using base protocol extensions) may help, especially if the data consumer's software is aware of such extensions.

2.2.2. Unauthorized Actions of Legitimate OPES Entities

According to the OPES architecture, the authorization is not tightly coupled with specific rules and procedures triggered by the rules. Even if a requirement to approve each particular rule and procedure was set, it looks at least impractical, if not impossible, to request such permission from the end user. Authorization granularity extends to transformation classes, but not to individual rules or transformations. The actual rules and triggered procedures may (maliciously or due to a programming error) perform actions that they are not authorized for.

2.2.3. Unwanted Content Transformations

An authorized OPES service may perform actions that do not adhere to the expectations of the party that gave the authorization for the service. Examples may include ad flooding by a local ad insertion service or use of inappropriate policy by a content filtering service.

On the other hand, an OPES entity acting on behalf of one party may perform transformations that another party deems inappropriate. Examples may include replacing ads initially inserted by the content provider or applying filtering transformations that change the meaning of the text.

2.2.4. Corrupted Content

The OPES system may deliver outdated or otherwise distorted information due to programming problems or as a result of malicious attacks. For example, a compromised server, instead of performing an OPES service, may inject bogus content. Such an action may be an act of cyber-vandalism (including virus injection) or intentional distribution of misleading information (such as manipulations with financial data).

A compromised OPES server or malicious entity in the data flow may introduce changes specifically intended to cause improper actions in the OPES server or callout server. These changes may be in the message body, headers, or both. This type of threat is discussed in more detail below.

2.2.5. Threats to Message Structure Integrity

An OPES server may add, remove, or delete certain headers in a request and/or response message (for example, to implement additional privacy protection or assist in content filtering). Such changes may violate end-to-end integrity requirements or defeat services that use information provided in such headers (for example, some local filtering services or reference-based services).

2.2.6. Granularity of Protection

OPES services have implicit permission to modify content. However, the permissions generally apply only to portions of the content, for example, URL's between particular HTML tags, text in headlines, or URL's matching particular patterns. In order to express such policies, one must be able to refer to portions of messages and to detect modifications to message parts.

Because there is currently very little support for policies that are expressed in terms of message parts, it will be difficult to attribute any particular modification to a particular OPES processor, or to automatically detect policy violations.

A fine-grained policy language should be devised, and it could be enforced using digital signatures. This would avoid the problems inherent in hop-by-hop data integrity measures (see next section).

2.2.7. Risks of Hop-by-Hop Protection

Generally, OPES services cannot be applied to data protected with end-to-end encryption methods because the decryption key cannot be shared with OPES processors without compromising the intended confidentiality of the data. This means that if the endpoint policies permit OPES services, the data must either be transmitted without confidentiality protections or an alternative model to end-to-end encryption must be developed, one in which the confidentiality is guaranteed hop-by-hop. Extending the end-to-end encryption model is out of scope of this work.

OPES services that modify data are incompatible with end-to-end integrity protection methods, and this work will not attempt to define hop-by-hop integrity protection methods.

2.2.8. Threats to Integrity of Complex Data

The OPES system may violate data integrity by applying inconsistent transformations to interrelated data objects or references within the data object. Problems may range from a broken reference structure

(modified/missing targets, references to wrong locations or missing documents) to deliberate replacement/deletion/insertion of links that violate intentions of the content provider.

2.2.9. Denial of Service (DoS)

The data consumer application may not be able to access data if the OPES system fails for any reason.

A malicious or malfunctioning node may be able to block all traffic. The data traffic destined for the OPES processor (or callout server) may not be able to use the services of the OPES device. The DoS may be achieved by preventing the data traffic from reaching the processor or the callout server.

2.2.10. Tracing and Notification Information

Inadequate or vulnerable implementation of the tracing and notification mechanisms may defeat safeguards built into the OPES architecture.

Tracing and notification facilities may become a target of malicious attack. Such an attack may create problems in discovering and stopping other attacks.

The absence of a standard for tracing and notification information may present an additional problem. This information is produced and consumed by the independent entities (OPES servers/user agents/content provider facilities). This calls for a set of standards related to each base protocol in use.

2.2.11. Unauthenticated Communication in OPES Flow

There are risks and threats that could arise from unauthenticated communication between the OPES server and callout servers. Lack of use of strong authentication between OPES processors and callout servers may open security holes whereby DoS and other types of attacks (see sections [2 and 3]) can be performed.

3. Threats to Out-of-Band Data

The OPES architecture separates a data flow from a control information flow (loading rulesets, trust establishment, tracing, policy propagation, etc.). There are certain requirements set for the latter, but no specific mechanism is prescribed. This gives more flexibility for implementations, but creates more burden for implementers and potential customers to ensure that each specific

implementation meets all requirements for data security, entity authentication, and action authorization.

In addition to performing correct actions on the OPES data flow, any OPES implementation has to provide an adequate mechanism to satisfy requirements for out-of-band data and signaling information integrity.

Whatever the specific mechanism may be, it inevitably becomes subject to multiple security threats and possible attacks. The way the threats and attacks may be realized depends on implementation specifics but the resulting harm generally falls into two categories: threats to OPES data flow and threats to data integrity.

The specific threats are:

3.1. Threats that Endanger the OPES Data Flow

Any weakness in the implementation of a security, authentication, or authorization mechanism may open the door to the attacks described in section 2.

An OPES system implementation should address all these threats and prove its robustness and ability to withstand malicious attacks or networking and programming problems.

3.2. Inaccurate Accounting Information

Collecting and reporting accurate accounting data may be vital when OPES servers are used to extend a business model of a content provider, service provider, or as a basis for third party service. The ability to collect and process accounting data is an important part of OPES' system functionality. This functionality may be challenged by distortion or destruction of base accounting data (usually logs), processed accounting data, accounting parameters, and reporting configuration.

As a result a data consumer may be inappropriately charged for viewing content that was not successfully delivered, or a content provider or independent OPES services provider may not be compensated for the services performed.

The OPES system may use accounting information to distribute resources between different consumers or limit resource usage by a specific consumer. In this case an attack on the accounting system (by distortion of data or issuing false configuration commands) may result in incorrect resource management and DoS by artificial resource starvation.

3.3. OPES Service Request Repudiation

An entity (producer or consumer) might make an authorized request and later claim that it did not make that request. As a result, an OPES entity may be held liable for unauthorized changes to the data flow, or will be unable to receive compensation for a service.

There should be a clear request that this service is required and there should be a clear course of action on behalf of all parties. This action should have a request, an action, a non-repudiable means of verifying the request, and a means of specifying the effect of the action.

3.4. Inconsistent Privacy Policy

The OPES entities may have privacy policies that are not consistent with the data consumer application or content provider application.

Privacy related problems may be further complicated if OPES entities, content providers, and end users belong to different jurisdictions with different requirements and different levels of legal protection. As a result, the end user may not be aware that he or she does not have the expected legal protection. The content provider may be exposed to legal risks due to a failure to comply with regulations of which he is not even aware.

3.5. Exposure of Privacy Preferences

The OPES system may inadvertently or maliciously expose end user privacy settings and requirements.

3.6. Exposure of Security Settings

There are risks that the OPES system may expose end user security settings when handling the request and responses. The user data must be handled as sensitive system information and protected against accidental and deliberate disclosure.

3.7. Improper Enforcement of Privacy and Security Policy

OPES entities are part of the content distribution system and as such take on certain obligations to support security and privacy policies mandated by the content producer and/or end user. However there is a danger that these policies are not properly implemented and enforced. The data consumer application may not be aware that its protections are no longer in effect.

There is also the possibility of security and privacy leaks due to the accidental misconfiguration or, due to misunderstanding what rules are in effect for a particular user or request.

Privacy and security related parts of the systems can be targeted by malicious attacks and the ability to withstand such attacks is of paramount importance.

3.8. DoS Attacks

DoS attacks can be of various types. One type of DoS attack takes effect by overloading the client. For example, an intruder can direct an OPES processor to issue numerous responses to a client. There is also additional DoS risk from a rule misconfiguration that would have the OPES processor ignore a data consumer application.

4. Security Considerations

This document discusses multiple security and privacy issues related to the OPES services.

5. References

5.1. Normative References

- [1] Barbir, A., Penno, R., Chen, R., Hofmann, M., and H. Orman, "An Architecture for Open Pluggable Edge Services (OPES)", RFC 3835, August 2004.
- [2] Barbir, A., Burger, E., Chen, R., McHenry, S., Orman, H., and R. Penno, "OPES Use Cases and Deployment Scenarios", RFC 3752, April 2004.
- [3] Barbir, A., Batuner, O., Beck, A., Chan, T., and H. Orman, "Policy, Authorization, and Enforcement Requirements of Open Pluggable Edge Services (OPES)", RFC 3838, August 2004.

5.2. Informative References

- [4] Floyd, S. and L. Daigle, "IAB Architectural and Policy Considerations for Open Pluggable Edge Services", RFC 3238, January 2002.

6. Acknowledgements

Many thanks to T. Chan (Nokia) and A. Beck (Lucent).

7. Authors' Addresses

Abbie Barbir
Nortel Networks
3500 Carling Avenue
Nepean, Ontario K2H 8E9
Canada

Phone: +1 613 763 5229
EMail: abbieb@nortelnetworks.com

Oskar Batuner
Independent consultant

EMail: batuner@attbi.com

Bindignavile Srinivas
Nokia
5 Wayside Road
Burlington, MA 01803
USA

EMail: bindignavile.srinivas@nokia.com

Markus Hofmann
Bell Labs/Lucent Technologies
Room 4F-513
101 Crawfords Corner Road
Holmdel, NJ 07733
US

Phone: +1 732 332 5983
EMail: hofmann@bell-labs.com

Hilarie Orman
Purple Streak Development

EMail: ho@alum.mit.edu

8. Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.