

Network Working Group
Request for Comments: 3914
Category: Informational

A. Barbir
Nortel Networks
A. Rousskov
The Measurement Factory
October 2004

Open Pluggable Edge Services (OPES) Treatment of
IAB Considerations

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

IETF Internet Architecture Board (IAB) expressed nine architecture-level considerations for the Open Pluggable Edge Services (OPES) framework. This document describes how OPES addresses those considerations.

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	Consideration (2.1) 'One-party consent'	3
4.	Consideration (2.2) 'IP-layer communications'	4
5.	Notification Considerations	5
5.1.	Notification versus trace.	6
5.2.	An example of an OPES trace for HTTP	8
5.3.	Consideration (3.1) 'Notification'	9
5.4.	Consideration (3.2) 'Notification'	10
6.	Consideration (3.3) 'Non-blocking'	10
7.	Consideration (4.1) 'URI resolution'	11
8.	Consideration (4.2) 'Reference validity'	11
9.	Consideration (4.3) 'Addressing extensions'	12
10.	Consideration (5.1) 'Privacy'	12
11.	Consideration 'Encryption'	12
12.	Security Considerations	13
13.	Compliance	13
14.	References	14
14.1.	Normative References	14
14.2.	Informative References	14
	Authors' Addresses	15
	Full Copyright Statement	16

1. Introduction

The Open Pluggable Edge Services (OPES) architecture [RFC3835], 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.

In the process of chartering OPES, the IAB made recommendations on issues that OPES solutions should be required to address. These recommendations were formulated in the form of a specific IAB considerations document [RFC3238]. In that document, IAB emphasized that its considerations did not recommend specific solutions and did not mandate specific functional requirements. Addressing an IAB consideration may involve showing appropriate protocol mechanisms or demonstrating that the issue does not apply. Addressing a consideration does not necessarily mean supporting technology implied by the consideration wording.

The primary goal of this document is to show that all formal IAB recommendations are addressed by OPES, to the extent that those considerations can be addressed by an IETF working group. The limitations of OPES working group to address certain aspects of IAB considerations are also explicitly documented.

IAB considerations document [RFC3238] contains many informal recommendations. For example, while the IAB informally requires OPES architecture to "protect end-to-end data integrity by supporting end-host detection and response to inappropriate behavior by OPES intermediaries", the IAB has chosen to formalize these requirements via a set of more specific recommendations, such as Notification considerations addressed in Section 5.3 and Section 5.4 below. OPES framework addresses informal IAB recommendations by addressing corresponding formal considerations.

There are nine formal IAB considerations [RFC3238] that OPES has to address. In the core of this document are the corresponding nine "Consideration" sections. For each IAB consideration, its section contains general discussion as well as references to specific OPES mechanisms relevant to the consideration.

2. Terminology

This document does not introduce any new terminology but uses terminology from other OPES documents.

3. Consideration (2.1) 'One-party consent'

"An OPES framework standardized in the IETF must require that the use of any OPES service be explicitly authorized by one of the application-layer end-hosts (that is, either the content provider or the client)." [RFC3238]

OPES architecture requires that "OPES processors MUST be consented to by either the data consumer or data provider application" [RFC3835]. While this requirement directly satisfies IAB concern, no requirement alone can prevent consent-less introduction of OPES processors. In other words, OPES framework requires one-party consent but cannot guarantee it in the presence of in-compliant OPES entities.

In [RFC3897], the OPES architecture enables concerned parties to detect unwanted OPES processors by examining OPES traces. While the use of traces in OPES is mandatory, a tracing mechanism on its own cannot detect processors that are in violation of OPES specifications. Examples include OPES processors operating in stealth mode. However, the OPES architecture allows the use of content signature to verify the authenticity of performed

adaptations. Content signatures is a strong but expensive mechanism that can detect any modifications of signed content provided that the content provider is willing to sign the data and that the client is willing to either check the signature or relay received content to the content provider for signature verification.

OPES entities may copy or otherwise access content without modifying it. Such access cannot be detected using content signatures. Thus, "passive" OPES entities can operate on signed content without the data consumer or provider consent. If content privacy is a concern, then content encryption can be used. A passive processor is no different from any intermediary operating outside of OPES framework. No OPES mechanism (existing or foreseeable) can prevent non-modifying access to content.

In summary, the one-party consent is satisfied by including the corresponding requirement in the IAB architecture document. That requirement alone cannot stop in-compliant OPES entities to perform consent-less adaptations, but OPES framework allows for various means of detecting and/or preventing such adaptations. These means typically introduce overheads and require some level of producer-consumer cooperation.

4. Consideration (2.2) 'IP-layer communications'

"For an OPES framework standardized in the IETF, the OPES intermediary must be explicitly addressed at the IP layer by the end user" [RFC3238].

The OPES architecture requires that "OPES processors MUST be addressable at the IP layer by the end user (data consumer application)" [RFC3835]. The IAB and the architecture documents mention an important exception: addressing the first OPES processor in a chain of processors is sufficient. That is, a chain of OPES processors is viewed as a single OPES "system" at the address of the first chain element.

The notion of a chain is not strictly defined by IAB. For the purpose of addressing this consideration, a group of OPES processors working on a given application transaction is considered. Such a group would necessarily form a single processing chain, with a single "exit" OPES processor (i.e., the processor that adapted the given message last). The OPES architecture essentially requires that last OPES processor to be explicitly addressable at the IP layer by the data consumer application. The chain formation, including its exit point may depend on an application message and other dynamic factors such as time of the day or system load.

5.1. Notification versus trace

Before specific considerations are discussed, the relationship between IAB notifications and OPES tracing has to be explained. OPES framework concentrates on tracing rather than notification. The OPES Communications specification [RFC3897] defines "OPES trace" as application message information about OPES entities that adapted the message. Thus, OPES trace follows the application message it traces. The trace is for the recipient of the application message. Traces are implemented as extensions of application protocols being adapted and traced.

As opposed to an OPES trace, provider notification (as implied by IAB) notifies the sender of the application message rather than the recipient. Thus, notifications propagate in the opposite direction of traces. Supporting notifications directly would require a new protocol. Figure 2 illustrates the differences between a trace and notification from a single application message point of view.




Figure 2

Since notifications cannot be piggy-backed to application messages, they create new messages and may double the number of messages the sender has to process. The number of messages that need to be proceed is larger if several intermediaries on the message path generate notifications. Associating notifications with application messages may require duplicating application message information in notifications and may require maintaining a sender state until notification is received. These actions increase the performance overhead of notifications.

The level of available details in notifications versus provider interest in supporting notification is another concern. Experience shows that content providers often require very detailed information about user actions to be interested in notifications at all. For example, Hit Metering protocol [RFC2227] has been designed to supply content providers with proxy cache hit counts, in an effort to reduce cache busting behavior which was caused by content providers desire to get accurate site "access counts". However, the Hit Metering protocol is currently not widely deployed because the protocol does not supply content providers with information such as client IP addresses, browser versions, or cookies.

Hit Metering experience is relevant because Hit Metering protocol was designed to do for HTTP caching intermediaries what OPES notifications are meant to do for OPES intermediaries. Performance requirements call for state reduction via aggregation of notifications while provider preferences call for state preservation or duplication. Achieving the right balance when two sides belong to different organizations and have different optimization priorities may be impossible.

Thus, instead of explicitly supporting notifications at the protocol level, OPES concentrates on tracing facilities. In essence, OPES supports notifications indirectly, using tracing facilities. In other words, the IAB choice of "Notification" label is interpreted as "Notification assistance" (i.e., making notifications meaningful) and is not interpreted as a "Notification protocol".

The above concerns call for making notification optional. The OPES architecture allows for an efficient and meaningful notification protocol to be implemented in certain OPES environments. For example, an OPES callout server attached to a gateway or firewall may scan outgoing traffic for signs of worm or virus activity and notify a local Intrusion Detection System (IDS) of potentially compromised hosts (e.g., servers or client PCs) inside the network. Such notifications may use OPES tracing information to pinpoint the infected host (which could be another OPES entity). In this example, notifications are essentially sent back to the content producer (the local network) and use OPES tracing to supply details.

Another environment where efficient and meaningful notification using OPES tracing is possible are Content Delivery Networks (CDNs). A CDN node may use multiple content adaptation services to customize generic content supplied by the content producer (a web site). For example, a callout service may insert advertisements for client-local events. The CDN node itself may not understand specifics of the ad insertion algorithm implemented at callout servers. However, the node may use information in the OPES trace (e.g., coming from the callout service) to notify the content producer. Such notifications may be about the number of certain advertisements inserted (i.e., the number of "impressions" delivered to the customer) or even the number of ad "clicks" the customer made (e.g., if the node hosts content linked from the ads). Callout services doing ad insertion may lack details (e.g., a customer ID/address or a web server authentication token) to contact the content producer directly in this case. Thus, OPES trace produced by an OPES service becomes essential in enabling meaningful notifications that the CDN node sends to the content producer.

5.2. An example of an OPES trace for HTTP

The example below illustrates adaptations done to HTTP request at an OPES processor operated by the client ISP. Both original (as sent by an end user) and adapted (as received by the origin web server) requests are shown. The primary adaptation is the modification of HTTP "Accept" header. The secondary adaptation is the addition of an OPES-System HTTP extension header [I-D.ietf-opes-http].

```
GET /pub/WWW/ HTTP/1.1
Host: www.w3.org
Accept: text/plain
```

Figure 3

... may be adapted by an ISP OPES system to become:

```
GET /pub/WWW/ HTTP/1.1
Host: www.w3.org
Accept: text/plain; q=0.5, text/html, text/x-dvi; q=0.8
OPES-System: http://www.isp-example.com/opes/?client-hash=1234567
```

Figure 4

The example below illustrates adaptations done to HTTP response at an OPES intermediary operated by a Content Distribution Network (CDN). Both original (as sent by the origin web server) and adapted (as received by the end user) responses are shown. The primary adaptation is the conversion from HTML markup to plain text. The secondary adaptation is the addition of an OPES-System HTTP extension header.

```
HTTP/1.1 200 OK
Content-Length: 12345
Content-Encoding: text/html

<html><head><h1>Available Documenta...
```

Figure 5

... may be adapted by a CDN OPES system to become:

```
HTTP/1.1 200 OK
Content-Length: 2345
Content-Encoding: text/plain
OPES-System: http://www.cdn-example.com/opes/?site=7654&svc=h2t

AVAILABLE DOCUMENTA...
```

Figure 6

In the above examples, OPES-System header values contain URIs that may point to OPES-specific documents such as description of the OPES operator and its privacy policy. Those documents may be parameterized to allow for customizations specific to the transaction being traced (e.g., client or even transaction identifier may be used to provide more information about performed adaptations). An OPES-Via header may be used to provide a more detailed trace of specific OPES entities within an OPES System that adapted the message. Traced OPES URIs may be later used to request OPES bypass [RFC3897].

5.3. Consideration (3.1) 'Notification'

"The overall OPES framework needs to assist content providers in detecting and responding to client-centric actions by OPES intermediaries that are deemed inappropriate by the content provider" [RFC3238].

OPES tracing mechanisms assist content providers in detecting client-centric actions by OPES intermediaries. Specifically, a compliant OPES intermediary or system notifies a content provider of its presence by including its tracing information in the application protocol requests. An OPES system MUST leave its trace [RFC3897]. Detection assistance has its limitations. Some OPES intermediaries may work exclusively on responses and may not have a chance to trace the request. Moreover, some application protocols may not have explicit requests (e.g., a content push service).

OPES tracing mechanisms assist content providers in responding to client-centric actions by OPES intermediaries. Specifically, OPES traces MUST include identification of OPES systems and SHOULD include a list of adaptation actions performed on provider's content. This tracing information may be included in the application request. Usually, however, this information will be included in the application response, an adapted version of which does not reach the content provider. If OPES end points cooperate, then notification can be assisted with traces. Content providers that suspect or experience difficulties can do any of the following:

- o Check whether requests they receive pass through OPES intermediaries. Presence of OPES tracing info will determine that. This check is only possible for request/response protocols. For other protocols (e.g., broadcast or push), the provider would have to assume that OPES intermediaries are involved until proven otherwise.

- o If OPES intermediaries are suspected, request OPES traces from potentially affected user(s). The trace will be a part of the application message received by the user software. If involved parties cooperate, the provider(s) may have access to all the needed information. Certainly, lack of cooperation may hinder access to tracing information. To encourage cooperation, data providers might be able to deny service to uncooperative users.
- o Some traces may indicate that more information is available by accessing certain resources on the specified OPES intermediary or elsewhere. Content providers may query for more information in this case.
- o If everything else fails, providers can enforce no-adaptation policy using appropriate OPES bypass mechanisms and/or end-to-end encryption mechanisms.

OPES detection and response assistance is limited to application protocols with support for tracing extensions. For example, HTTP [RFC2616] has such support while DNS over UDP does not.

5.4. Consideration (3.2) 'Notification'

"The overall OPES framework should assist end users in detecting the behavior of OPES intermediaries, potentially allowing them to identify imperfect or compromised intermediaries" [RFC3238].

OPES tracing mechanisms assist end users in detecting OPES intermediaries. Specifically, a compliant OPES intermediary or system notifies an end user of its presence by including its tracing information in the application protocol messages sent to the client. An OPES system MUST leave its trace [RFC3897]. However, detection assistance has its limitations. Some OPES systems may work exclusively on requests and may not have a chance to trace the response. Moreover, some application protocols may not have explicit responses (e.g., event logging service).

OPES detection assistance is limited to application protocols with support for tracing extensions. For example, HTTP [RFC2616] has such support while DNS over UDP does not.

6. Consideration (3.3) 'Non-blocking'

"If there exists a "non-OPES" version of content available from the content provider, the OPES architecture must not prevent users from retrieving this "non-OPES" version from the content provider" [RFC3238].

"OPES entities MUST support a bypass feature" [RFC3897]. If an application message includes bypass instructions and an OPES intermediary is not configured to ignore them, the matching OPES intermediary will not process the message. An OPES intermediary may be configured to ignore bypass instructions only if no non-OPES version of content is available. Bypass may generate content errors since some OPES services may be essential but may not be configured as such.

Bypass support has limitations similar to the two notification-related considerations above.

7. Consideration (4.1) 'URI resolution'

"OPES documentation must be clear in describing these services as being applied to the result of URI resolution, not as URI resolution itself" [RFC3238].

"OPES Scenarios and Use Cases" specification [RFC3752] documents content adaptations that are in scope of the OPES framework. Scenarios include content adaptation of requests and responses. These documented adaptations do not include URI resolution. In some environments, it is technically possible to use documented OPES mechanisms to resolve URIs (and other kinds of identifiers or addresses). The OPES framework cannot effectively prevent any specific kind of adaptation.

For example, a CDN node may substitute domain names in URLs with CDN-chosen IP addresses, essentially performing a URI resolution on behalf of the content producer (i.e., the web site owner). An OPES callout service running on a user PC may rewrite all HTML-embedded advertisement URLs to point to a user-specified local image, essentially performing a URI redirection on behalf of the content consumer (i.e., the end user). Such URI manipulations are outside of the OPES framework scope, but cannot be effectively eliminated from the real world.

8. Consideration (4.2) 'Reference validity'

"All proposed services must define their impact on inter- and intra-document reference validity" [RFC3238].

The OPES framework does not propose adaptation services. However, OPES tracing requirements include identification of OPES intermediaries and services (for details, see "Notification" consideration sections in this document). It is required that

provided identification can be used to locate information about the OPES intermediaries, including the description of impact on reference validity [RFC3897].

9. Consideration (4.3) 'Addressing extensions'

"Any services that cannot be achieved while respecting the above two considerations may be reviewed as potential requirements for Internet application addressing architecture extensions, but must not be undertaken as ad hoc fixes" [RFC3238].

OPES framework does not contain ad hoc fixes. This document in combination with and other OPES documents should be sufficient to inform service creators of IAB considerations. If a service does URI resolution or silently affects document reference validity, the authors are requested to review service impact on Internet application addressing architecture and work within IETF on potential extension requirements. Such actions would be outside of the current OPES framework.

10. Consideration (5.1) 'Privacy'

"The overall OPES framework must provide for mechanisms for end users to determine the privacy policies of OPES intermediaries" [RFC3238].

OPES tracing mechanisms allow end users to identify OPES intermediaries (for details, see "Notification" consideration sections in this document). It is required that provided identification can be used to locate information about the OPES intermediaries, including their privacy policies.

The term "privacy policy" is not defined in this context (by IAB or OPES working group). OPES tracing mechanisms allow end users and content providers to identify an OPES system and/or intermediaries. It is believed that once an OPES system is identified, it would be possible to locate relevant information about that system, including information relevant to requesters perception of privacy policy or reference validity.

11. Consideration 'Encryption'

"If OPES is chartered, the OPES working group will also have to explicitly decide and document whether the OPES architecture must be compatible with the use of end-to-end encryption by one or more ends of an OPES-involved session. If OPES was compatible with end-to-end encryption, this would effectively ensure that OPES boxes would be

restricted to ones that are known, trusted, explicitly addressed at the IP layer, and authorized (by the provision of decryption keys) by at least one of the ends" [RFC3238].

The above quoted requirement was not explicitly listed as on of the IAB considerations, but still needs to be addressed. The context of the quote implies that the phrase "end-to-end encryption" refers to encryption along all links of the end-to-end path, with the OPES intermediaries as encrypting/decrypting participants or hops (e.g., encryption between the provider and the OPES intermediaries, and between the OPES intermediaries and the client).

Since OPES processors are regular hops on the application protocol path, OPES architecture allows for such encryption, provided the application protocol being adapted supports it. Hop-by-hop encryption would do little good for the overall application message path protection if callout services have to receive unencrypted content. To allow for complete link encryption coverage, OPES callout protocol (OCP) supports encryption of OCP connections between an OPES processor and a callout server via optional (negotiated) transport encryption mechanisms [I-D.ietf-opes-ocp-core].

For example, TLS encryption [RFC2817] can be used among HTTP hops (some of which could be OPES processors) and between each OPES processor and a callout server.

12. Security Considerations

This document does not define any mechanisms that may be subject to security considerations. This document scope is to address specific IAB considerations. Security of OPES mechanisms are discussed in Security Considerations sections of the corresponding OPES framework documents.

For example, OPES tracing mechanisms assist content providers and consumers in protecting content integrity and confidentiality by requiring OPES intermediaries to disclose their presence. Security of the tracing mechanism is discussed in the Security Considerations section of [RFC3897].

13. Compliance

This document may be perceived as a proof of OPES compliance with IAB implied recommendations. However, this document does not introduce any compliance subjects. Compliance of OPES implementations is defined in other OPES documents discussed above.

14. References

14.1. Normative References

- [RFC3238] Floyd, S. and L. Daigle, "IAB Architectural and Policy Considerations for Open Pluggable Edge Services", RFC 3238, January 2002.
- [RFC3752] Barbir, A., Burger, E., Chen, R., McHenry, S., Orman, H. and R. Penno, "Open Pluggable Edge Services (OPES) Use Cases and Deployment Scenarios", RFC 3752, April 2004.
- [RFC3835] Barbir, A., Penno, R., Chen, R., Hofmann, M., and H. Orman, "An Architecture for Open Pluggable Edge Services (OPES)", RFC 3835, August 2004.
- [RFC3897] Barbir, A., "Open Pluggable Edge Services (OPES) Entities and End Points Communication", RFC 3897, September 2004.

14.2. Informative References

- [RFC2227] Mogul, J. and P. Leach, "Simple Hit-Metering and Usage-Limiting for HTTP", RFC 2227, October 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P. and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC2817] Khare, R. and S. Lawrence, "Upgrading to TLS Within HTTP/1.1", RFC 2817, May 2000.
- [I-D.ietf-opes-http] Rousskov, A. and M. Stecher, "HTTP adaptation with OPES", Work in Progress, October 2003.

[I-D.ietf-opes-ocp-core]

Rousskov, A., "OPES Callout Protocol
Core", Work in Progress, November 2003.

Authors' Addresses

Abbie Barbir
Nortel Networks
3500 Carling Avenue
Nepean, Ontario
CA

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

Alex Rousskov
The Measurement Factory

EMail: rousskov@measurement-factory.com
URI: <http://www.measurement-factory.com/>

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/S HE 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 IETF's procedures with respect to rights in IETF 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.