Fendix Media Ltd

Segmented Advertising System

Process and Data Capture Overview

System Version 4.0.0

> Document Release Date: 16 June 2016

Introduction

The purpose of this document is to describe how the segmented advertising system delivers targeted adverts using Active Directory (AD) security groups on a Windows network. It also explains what information is transmitted to Fendix Media's servers and how this is used and stored.

Although primarily intended to answer commonly asked questions from IT teams, this guide will also be of interest to those with governance considerations regarding the transmission and use of an organisations' data.

Process overview

The following diagram provides a simplified schematic of the various stages of the segmentation process.

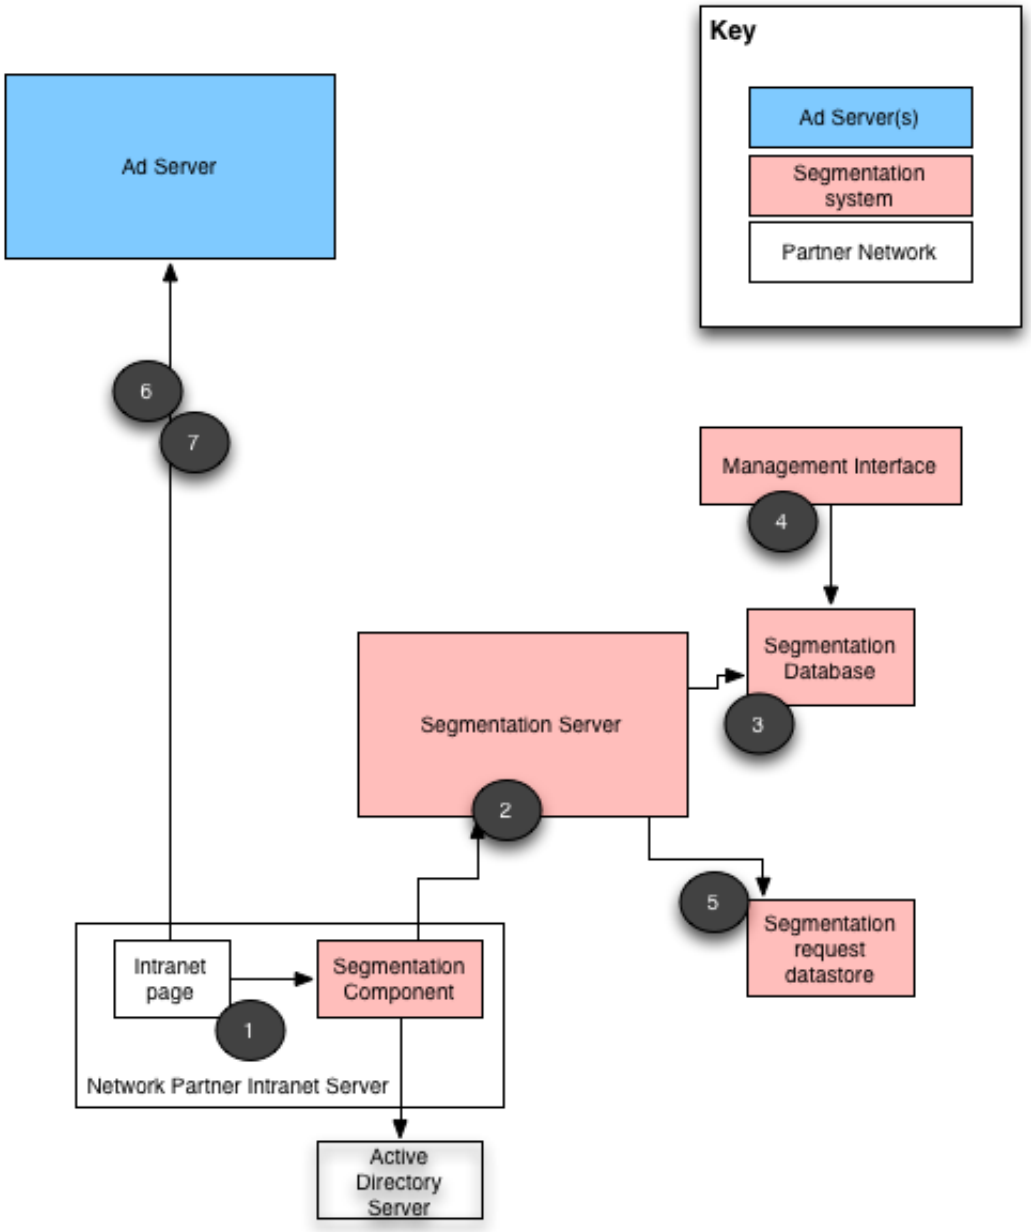

1. A small script tag embedded on the network partner's Intranet pages calls through to the Fendix Segmentation Component. This component is installed on a Microsoft web server (IIS) with Windows Authentication enabled - this authentication method allows the component to obtain the Windows username of the user viewing the Intranet page.

   The user name is then passed to the network AD server to request a list of AD groups of which the user is a member.
2. Once the list of groups has been obtained, it is encoded using Base64 encoding and transmitted to a web service on the Fendix segmentation server (code.fendixmedia.net), along with an MD5 hash of the username.
3. The segmentation server checks the list of supplied groups against a lookup table in the

segmentation database which matches AD group names to advertising keywords. This list is performed for each group to obtain an aggregate list of keywords pertinent to the user.

Once the list is complete, the segmentation server generates a small JavaScript tag specifying a call to the Fendix Ad Server containing the relevant keywords.

4. If a group name is encountered that is not already present in the segmentation database it is stored for later keywording using the Management interface, which allows designated Fendix staff to maintain keyword groupings for Network Partners.
5. For statistical purposes a record of each segmentation request is stored containing the following information:
   - Date/time of request
   - Network Partner unique reference number (URN)
   - User Hash
   - Keyword list
6. Once the relevant ad invocation code is passed back to the Segmentation Component, it is then returned to the Intranet page on the user's browser, where it is executed resulting in a call to the Fendix Ad Server to request an appropriate advert.
7. The advert code is returned to the user's browser where it is displayed.

Summary of data transfer

This section details which data is passed outside of the Network Partner network and what it is used for. Some explanatory notes are also provided to assist non technical readers who may not be familiar with specific terminology.

The following information is sent to the segmentation server:

Trust Unique code: This is the 4 character code which identifies each trust and is supplied by Fendix Media.

AD Group List: The full list of AD security groups that the user is a member of. These are used by the segmentation server to match to advertising keywords via a look up mechanism.

Group names and name-keyword mappings are stored in the main segmentation database. No user specific information is stored in this database.

MD5 Username Hash: A hash is a one way mathematical function that encodes a piece of data in a consistent way, but such that it is not possible to decode the message to view the original content. By using a hash of the username, it is possible for Fendix Media to gain insight into the number of unique users that match to each ad keyword whilst ensuring that it is not possible to identify actual people

Key datastores

There are three key datastores that form the Fendix ad-serving infrastructure, the segmentation database, the core ad server(s) and the analytics database. Each resides on separate servers with no linkages between them.

Segmentation database: Stores lists of AD groups for each trust along with keyword mappings.

No user specific data is stored, nor is individual ad request data stored in this database. Access to the segmentation database is allowed via the segmentation API servers (these run as a cluster of servers for availability and performance) and the Fendix Management Interface application server.

Analytics database: this is a document store which records a log of every segmented ad request, this log records the following data:

- Date/time of request
- Network Partner unique reference number (URN)
- User Hash
- Keyword list

From this data, analytical information such as the number of ad requests containing particular keywords and the number of unique individuals who match a given keyword can be deduced. This information is of interest to advertisers for assessing potential reach of an advertising campaign. It also allows Fendix Media to assess trends. None of the raw data is available to advertisers

Access to the analytics database is only possible from the segmentation servers (these only perform writes to the database), with read access only possible from within the Fendix media office network.

Ad Servers: The core advertising servers are where actual ad content is delivered from, these are run as a managed service by the system developers and function completely independently of the segmentation infrastructure.

No Trust data is ever communicated to the ad servers.

Administrative access

Administrative access to all servers in the Fendix segmentation infrastructure is by secure SSH session using public/private key access. Private keys are only available to Fendix system administrators and where applicable, access is only permissible from IP addresses within the Fendix network.

Information dissemination

Advertisers are able to select the broad specialism keywords that they wish to target their ads for serving to (e.g. doctor's, procurement, cardiac). They are able to choose these from a pre-defined list. For example, Macmillan Cancer Support recently ran a campaign to inform healthcare professionals of the support information available to patients. They wanted to target this campaign to nurses, cancer and palliative care. The only information that is supplied to advertisers is the number of ad impressions served against the chosen keywords (if they request this level of granularity, typically an advertiser is only given a report of the number of impressions served and the number of clicks generated)

At no point do advertisers have access to actual Active Directory group names, nor how they are mapped to specialism keywords. Advertisers are also not advised as to how many groups are mapped to a keyword, nor how many users have membership of a given group.