Responses to Frequently Asked Questions:

How do we protect anonymity of individuals whose job titles effectively identify them? (Just to give a couple of examples, Chief Executive, Head of Communications …. there are lots of others.)

We protect privacy by never storing the hashed user identifier along with the role or any other information that is used to allocate ad keywords. When a segmentation request is made we store the user hash details along with which ad keywords that were allocated to that request, not the AD group or ESR entry.

As an example, if the role, Chief executive was recognised, it would be assigned the keywords or 'directors and 'senior managers' - This is the information that we would store against the user hash. As many roles/job titles etc can be mapped to each keyword, it is not possible to map keywords back to a specific role.

How do we deal with the fact that staff are not being able to opt in (since this is for work purposes using work IT, is it actually any different from feeding screensavers to all machines? Or, could it give rise to a complaint to the ICO?)

Our current basis for processing data is that of legitimate interest - via the commercial relationship between the Trust and Fendix Media. It should be noted that we do not track any users activities, nor do we partake in any third party tracking programs. We also do not store any personally identifiable information and we do not share information with third parties.

Where is your ad server hosted and how secure is it? What precautions do you take to ensure security? (We need to address the risk of maladvertising)

All of our servers are hosted in the EU (Ireland for segmentation servers, Sweden for the ad servers – these are run for us as a managed service by the developers, who are responsible for all security and patching). For security and access controls we employ see the process document.

How do we protect our IT systems in the event of your server being hacked?

The simplest option is just to switch off the serving switch in the configuration of the trust side component – this disables the entire process. Our servers do not require any inbound access to the trust network, so it is no different to any other website or service.

If, for whatever reason, you have any issues with your ad-server, could it potentially cause issues for our IT system (eg causing our intranet to crash or go slow)

We recommend loading ads asynchronously using the third part post-scribe library, which loads ads in the background without delaying the rest of page load – that way in the event of an issue with ad serving (or as has proved much more likely, issue with the N3) it does not affect page load

How can we be sure your pseudonymysation is completely reliable?

We hash the username using Md5, SHA256 or SHA512 hashing algorithms – hashes are none reversible and given that cracking the hash would bring no benefits to an attacker it would be pointless to brute force attack.

If you are collecting information about pages visited or click-throughs in order to report back to advertisers, what guarantees can you give that you do not share or further process data?

The only information we share with advertisers is statistical information. We do not share tracking or personal data of any form – to do so would be a GDPR violation as we do not have opt ins to share personal info with any named third parties.

Can staff opt-out of the ads – you can do this on your personal devices you can ask that your information or preferences not be shared?

We never share information or preferences, nor do we track users. It is not currently possible to opt out of ads.

- Can you switch if off individually?

With the ESR based segmentation, targeted ads could be disabled on a per individual basis by removing the relevant individual from the ESR data file.

- Will it slow down the Intranet?

No, ads can be loaded asynchronously using an open source script library that as designed specifically for ad loading – this way the ads load in the background without causing the page to wait

- Do they regularly check links and remove links that could be potentially unsafe?

Yes, we only accept campaigns directly from advertising clients and trusted ad agencies that we have a direct relationship with. All campaigns are checked prior to release to the Trust and all campaigns are explicitly approved by the trust prior to delivery.

- What's the guarantee of the security of the sites going to?

See above - as we work directly with the agency or end client all click through destinations are known. As the destination sites are on the open internet, normal considerations regarding site safety apply. Landing pages are communicated to the Trust as part of the campaign approval process prior to delivery

- "share the list of the sites that need to be whitelisted to your IT teams during the installation phase" we won't just open a site – if its outside policy or unsafe we wouldn't allow

Our standard recommended whitelist is at:

https://www.fendixmedia.co.uk/segmentation/FMWhitelist.pdf

This covers the most common domains that ads are delivered from by partner agencies. If the trust chooses not to whitelist any of these domains it may impact eligibility for some campaigns and thus revenue.

- Where is this run from locally or cloud based?

There is a small component that resides locally on the trust network, the rest of the segmentation system is hosted by Fendix Media and our core ad servers are run for us as a manged service by the system developers.

- What disclaimers are there around content of the sites e.g. UHMB are not liable

The trust can place disclaimer text either above or below the banner. Alternatively a disclaimer message can be injected along with the banner which allows for an information message to be displayed if the user clicks an icon in the corner of the banner.

- AD accounts and security groups will need to be really up-to-date

With the new ESR data based approach, you will need to have forenames and surnames populated in your AD records. AD groups are not used in this approach (please see the process documentation for further details.

One of our governance queries is whether you could use the information you collect about our group memberships to identify and track individuals. For example who's a member of the <Group redacted>, <Group redacted>, and <Group redacted> groups? – that would only be <Name redacted>, the Head of IT. Could advertisers create 'Venn-diagram' relationships that would allow for direct targeting of pretty much every key person in this Trust based on just their group membership. Which means the anonymous advertising can actually be targeted at a named person, should the advertiser wish it?

Our segmented advertising implementation only uses active directory groups to identify broad areas of specialisms, e.g. nurses, midwives etc, we do not store usernames. When a request is made for a segmented advert, the user AD group membership list is sent to our segmentation server where a lookup is performed to allocate the relevant keywords. It is these keywords not the AD groups that are passed to the ad server for selecting an ad to display

Many AD groups may be mapped to a single keyword , e.g. groups Oncology_Nurses and Oncology_Consultants would both be mapped to the keyword 'cancer'. Similary, an AD group can be mapped to multiple keywords, e.g. Oncology_Nurse might be mapped to the 'cancer' and 'nurses' keyword.

As a result of this many to many mapping it is not possible to identify a single individual from the keyword list, since there may be multiple group combinations that would generate the same list.

Can I ask what input you get from would-be advertisers? And what information do you supply to them? Do they get to know group names?

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.

We're a bit concerned about the creation of new groups; if we were to, for example, create a new group to investigate the building of a new hospital or building, the group name would clearly say what it's for and that could lead to sensitive information leaving the organisation. I can see various building contactors being interested in a group called "new hospital build - <Location redacted>"

As described above, no AD group names or personally identifying data are ever shared with advertisers (or anyone else outside of Fendix Media, including other partners. Within Fendix, group data is also only accessible to IT administrators and specific personal who perform keyword mapping operations). In the example you gave, the most granular level of identification that would be given would be a keyword such as 'facilities' or 'estates' which would be indistinguishable from other groups which were mapped to the same keyword.

It should also be noted that the mapping of groups to keywords is a manual, managed operation and that any groups can be excluded from the mapping process by request.

We are constantly reviewing our mapping procedures to ensure the relevance of the results and also investigating ways of involving trust partners directly in the mapping process, to ensure both accuracy and to provide an additional layer of control over privacy concerns

What trust information is transmitted to Fendix media and how is it used?

The following information is transmitted to the segmentation servers:

- List of Active Directory security group membership - these are used to allocate keywords for ad targeting

- MD5 hash of username - this is used to enable us to gather statistics for the number of unique users within each keyword. It is not used to track individual users. An MD5 hash is a non human readable representation of the original username. In addition hashing is a one-way mathematical process that cannot be reversed to obtain the

original username. At no point do Fendix Media receive actual usernames.

The segmentation service accepts secure HTTPS requests.

This data is only held on our servers and access is not shared with anyone outside of Fendix Media, no trust data is transmitted outside of the Fendix server infrastructure. Advertising organisations are only supplied with broad statistical data for the number of impressions for targeting keywords. Access to data within Fendix is restricted to authorised individuals solely for the purpose of keywording.

Administrative access to our servers is via secure key based SSH access, the private keys of which are only available to Fendix system administrators.