FTICKS Logging

The Jisc NRPS keep a log of successful roams between organisation connected to them. However, about 90% of roams happen within Federations so this data is not visible. To gain a complete picture of roaming Federations need to send their own roaming information to a central point.

The mechanism for accomplishing this is Syslog, a standard (RFC 5424) approach for logging information to centralised repositories. Jisc uses ELK (Elasticsearch, Logstash and Kibana) to handle, process and display roaming data.

For doing a similar task within eduroam GEANT devised a standard syslog message body format and called it FTICKS. Their requirements were:

1. Allow receipt of statistics events in a decentralised manner (i.e. from arbitrary, but legitimate sources). 2. Support semantics of established eduroam statistics collection, in particular: a. Quantify number of authentications carried out, noting source country of user and country visited. b. Quantify number of roaming days (total number of distinct MAC addresses seen roaming on a given day). c. Separate actual user traffic from automatically generated probe traffic (monitoring traffic). 3. Be implementable by participants (FLR or SP/IdP servers) in a lightweight manner; ideally completely stateless for the participant. 4. Contain a reliable duplicate detection. 5. Require only the bare minimum of information about users to satisfy the quantification goals in requirement 2. 6. Enable participants to opt in to receive more detailed statistics than those stated in requirement 2 (at the expense of giving away more information). 7. Be extensible enough to allow for future adaptation if changes are made to the eduroam infrastructure. 8. Be independent of the server software used by participants.

The eduroam model is very flat with just NRPS \rightarrow ORPS \rightarrow IdP, but the govroam model goes a bit further with, potentially, multiple levels of Region Proxies e.g. NRPS \rightarrow RRPS \rightarrow RRPS \rightarrow ... \rightarrow ORPS \rightarrow IdP.

Each site keeps logs internally which is enough for identifying local issues and passing on details for audit purposes. JISC's NRPS can see all the roaming interactions between sites and aren't interested in local users using eduroam on their own site.

However, the govroam model means that there could (and are) users roaming between sites which have their own RRP and thus the traffic isn't seen by the JISC NRPS. People are more likely to roam to a nearby site as opposed to a more remote one so most of the roaming traffic will be seen only by the RRPS and much less by the NRPS. Thus JISC is missing out on a considerable amount of meta-data describing the roaming activity of govroam users.

There's a requirement to see if and how logging information could be passed from the RRPS to a central JISC logging site. The obvious answer would be to have the RRPS configured to send FTICKS (the standard logging format used by eduroam and govroam) to the same logging server receiving the NRPS logs.

The problems start with the RADIUS software. Whilst it's easy to configure software such as RADIATOR (used by eduroam NRPS) and FreeRADIUS (used by the govroam NRPS) to generate logs in arbitary formats and send them to arbitary destinations, other software, such as Microsoft NPS appears to have no ability to do this. However, there might be tools which take the logs and convert them to the right format to be sent out. A quick look suggests that it's possible but there might be costs associated with the software. There might also be licence issues with buying software to then give to

Last update: 2021/04/28 10:07

someone else to use. An alternative would be writing something in-house which would avoid both of these problems.

ELK (ElasticSearch, Logstash, Kibana) is a monitoring tool currently being used to process the NRPS logs and is likely to be the ultimate destination of any logs generated by the RRPS/ORPS so integration with this is vital. Logstash can take logs in a variety of formats, including Windows Event logs (from a local machine), syslog and any number of log formats as long as data can be pattern matched. Logstash feeds directly into ElasticSearch though, making it a fairly heavy duty solution. Along with this triumverate is 'filebeat', a light-weight piece of software that can gather logs and send them onto Logstash. It's less capable of processing data but is easier to install and configure. The responsibility then falls on Logstash to process the data and store it in ElasticSearch.

Filebeat is available for Linux and Windows so should cover most installs. It's capable of reading the data from file logs generated by NPS at least.

Discussions would be needed with various sites running various different type of RADIUS software to determine what's possible and acceptable. Security could be an issue at a policy level. The ELK software is all capable of communicating securely but sites may have security policies which wouldn't allow traffic to be sent directly from internal RADIUS servers offsite.

Log Attributes

If extracting the data from the RADIUS server and sending it on to a log processing system has issues that can be overcome then the next problem, less surmountable, is the actual data generated by the RADIUS servers. Most information is available to the IdP, which has visibility of the inner tunnel, but there is no expectation, or even desire, for JISC to be presented with this data. Conversely, a RADIUS proxy can only see the data in the outer tunnel but can add extra information, such as source IP, to the log data. This is what which JISC needs to see.

One of the key data is the Operator-Name. This is an attribute set in the RADIUS server configuration and identifies the site sending the request. Manchester University would have the Operator-Name set to '1manchester.ac.uk', for example. If the sending site doesn't set this then the proxy can log information which identifies the source. This is less desirable as it only identifies the previous hop. With eduroam this wouldn't be a problem due to the flat nature but is more significant with govroam due to the number of layers of the hierarchy.

Insisting that all sites set the Operator-Name is problematic because some RADIUS servers don't have the capability of adding attributes (Microsoft NPS for example), but should still proxy any existing such attribute.

In a perfect world the first proxy to see an authentication request should set its Operator-Name to a globally unique (based on DNS for example) name and then pass the request up the chain. If that happened then whichever proxy is at the 'top' of the chain will have a log of exactly which site generated the initial request. Since the request also contains the realm of the visitor there is then a nice mapping of a roam. If all sites were capable of generating, collecting and forwarding this data then JISC would have a complete set of data describing roaming activity between Govroam enabled sites.

We may have to accept that, until NPS is fixed/improved, that there are going to be big holes in the collected data. As for other systems, documentation will be needed covering the install and configuration of whatever software used to generate and pass on F-TICKS or equivalent.

From:

https://wiki.govroam.uk/dokuwiki/ - Govroam

Permanent link:

https://wiki.govroam.uk/dokuwiki/doku.php?id=public:fticks&rev=1619604434

Last update: 2021/04/28 10:07

