# High Level Architectural Design

## Summary

Govroam uses a hierarchical tree of RADIUS servers to proxy authentication requests and responses around the country. The purpose is to allow people with valid credentials on one site to roam to another site without having to be assigned new credentials.

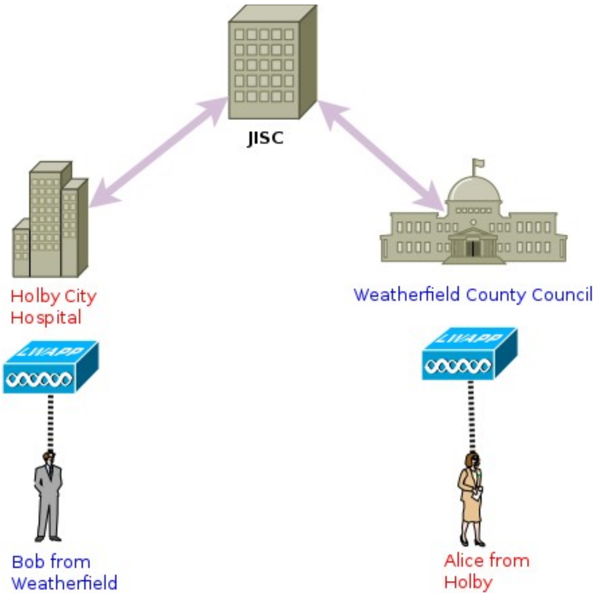

## Logical

The authentication process is:

1. A user connects their device to a specific SSID (govroam) when visiting another site.
2. An authentication request is sent from the Wireless system with a realm set. e.g. @holby.nhs.uk
3. This request is proxied, using the realm to 'route' to the request across the network of proxies.
4. The request finishes up at the site which can authenticate the user's credentials.
5. A response is sent back via the network of proxies.
6. The Wireless system uses the response to decide whether or not to let the user onto the data network.
7. Once on the network the user obtains an IP address and can access Internet as a visitor to the site.

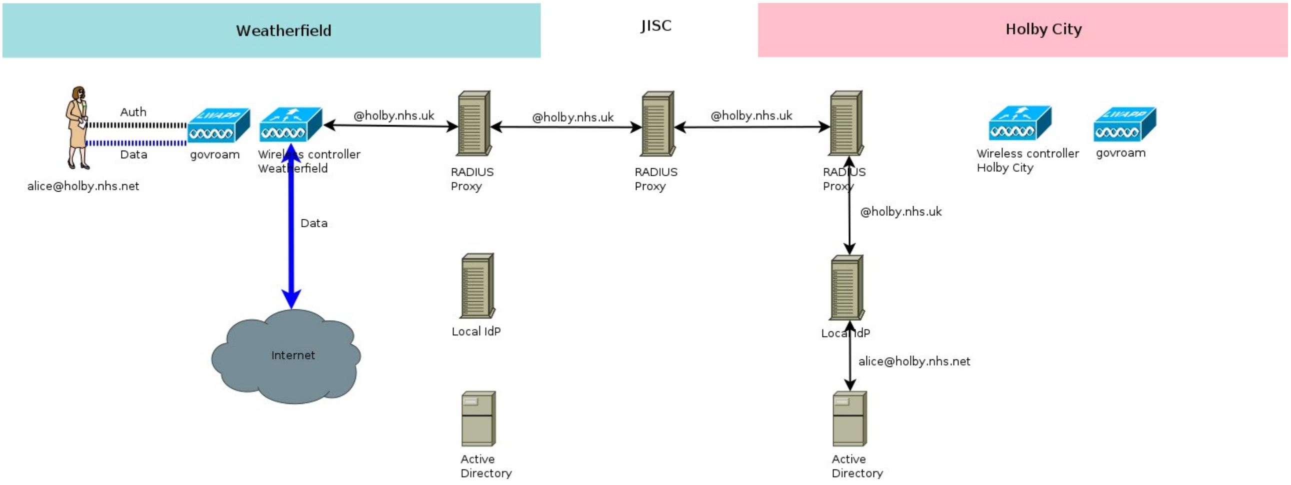

## Physical

Jisc runs the RADIUS servers (National RADIUS Proxy Servers, or NRPS) at the top of tree. These RADIUS servers are responsible for proxying between the sites whose RADIUS servers peer directly with them. These could be individual sites or federations.

Each site would have a wireless system, a repository of credentials, a RADIUS server to perform using these credentials (IdP), a RADIUS server to act as a proxy (ORPS/RRPS) and a connection to the Internet (with all the usual services such as DHCP and DNS).

At the most simple this could be one wireless AP, a server running Microsoft AD and NPS, with NPS configured to do both authentication and proxying, and an ADSL line.

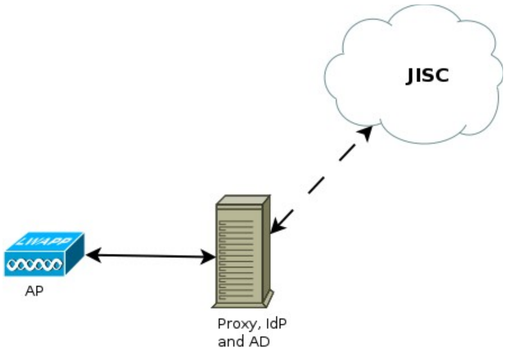

*Illustration 1: Simple Layout*

At the other end of the spectrum, it could be a few hundred sites spread across the UK each with their own enterprise wireless network, a number of AD servers serving multiple realms and groups of RADIUS servers, some acting as IdPs and some as proxies. There could be multiple levels of proxy depending on the logical, physical and security requirements of the sites.

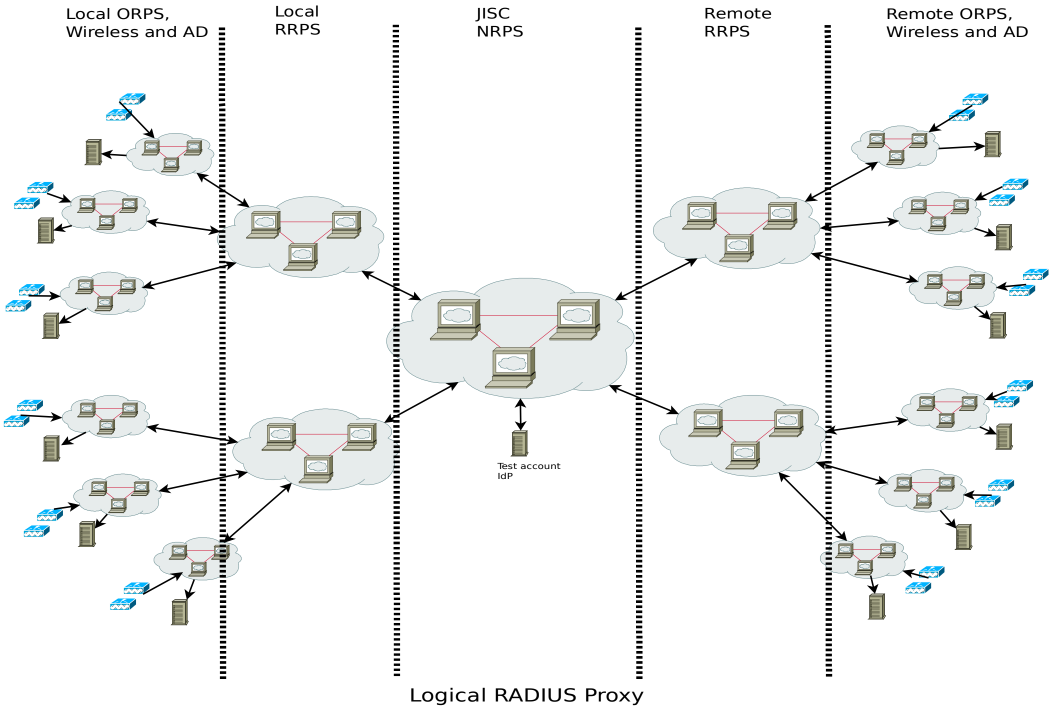

*Illustration 2: Complex Layout*

## Technical Details

The protocols and standards involved are known as 802.1x, EAP or WPA2-Enterprise. These terms are used a lot around systems like govroam. There is a lot of overlap between them so are often used interchangeably. A quick summary is that WPA2-Enterprise is a wireless standard which incorporates 802.1x to handle the authentication. In turn 802.1x defines how the EAP protocol is used. EAP (or Extensible Authentication Protocol) is a very flexible protocol for handling authentication requests. EAP itself simply encapsulates other protocols such as PAP and MSCHAPv2 which actually do the authentication.

These protocols are key to the way that govroam works.

The wireless SSID (govroam) has to be configured to use 802.1x as this is the foundation of the whole system.

The client has a built-in *supplicant* which is just a piece of software that handles the collection, and sending, of credentials. The initial transaction takes place between the supplicant and the wireless system. The supplicant has three key pieces of information: username, password and realm. The first two are obvious but the third is less so. The realm identifies the site and is normally the site's DNS name. It has to be a globally unique value so normally a DNS domain is used.

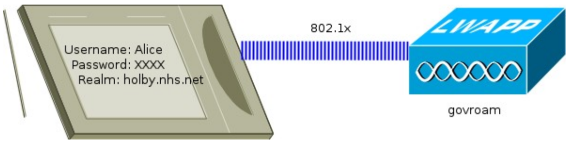

In turn the wireless system (often referred to the Network Access Server or NAS) generates a RADIUS request which is sent to a proxy server. In govroam nomenclature this is called an Organisational RADIUS Proxy Server (or ORPS for short).

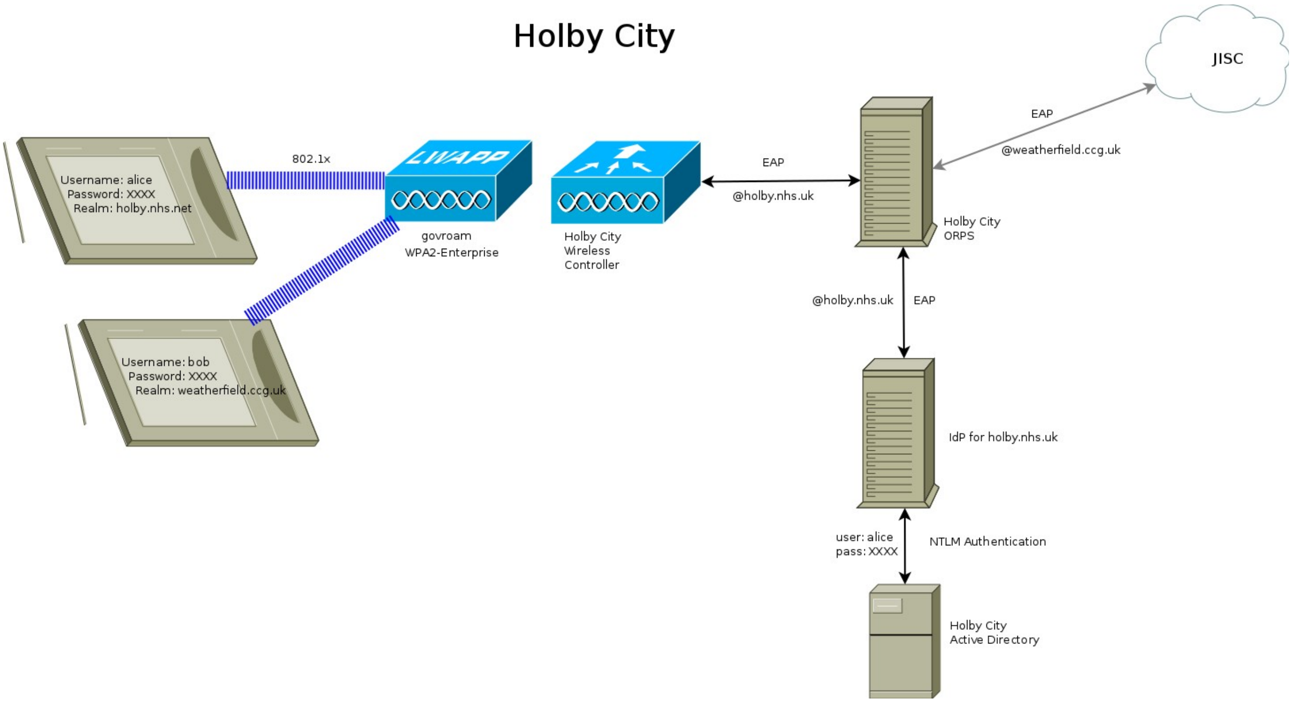

The ORPS uses the realm to decide where to send the authentication request. If it knows about the realm (i.e. it's local to the site) then it'll pass the request to the RADIUS IdP server associated with the realm. If it's unknown (in the case of a visitor to the site) then it'll pass the request up to the next level of proxies.

If the request arrives at the Jisc RADIUS proxies, the top of the tree, then it'll be passed down to which ever site is associated with that realm which, in turn, passes it down to another proxy until it arrives at an IdP and is authenticated (or not).

Now this has to work both ways. The above describes the process from the point of view of a visitor to your site but it's exactly the same if one of your users visits another site. The other site goes through the same process and the request eventually arrives at your site and proxied down to your IdP, authenticates and responds.

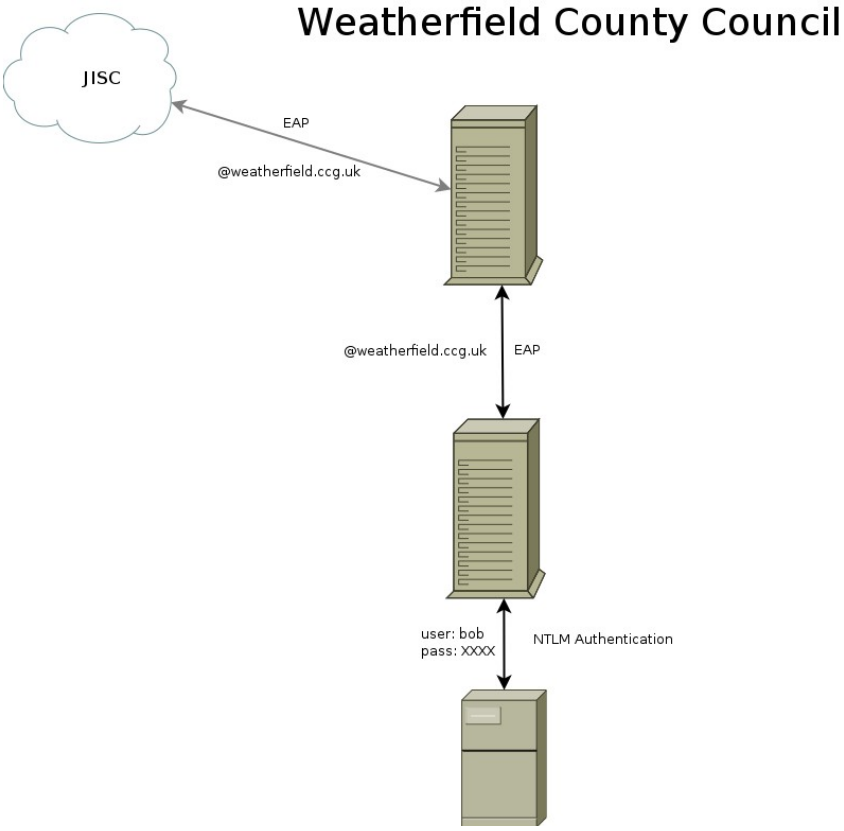

## More Technical Details

### Security

The above is a simplification of the process, as you might expect. You might be asking why would any one send their credentials across the Internet for anyone to see? If that were the case then this would be a bad idea. What actually happens is that the EAP protocol takes care of the security.

Logically EAP creates a tunnel through which the credentials are sent. The outer part of the tunnel uses SSL certificates (at one or both ends) to encrypt the contents of the inner tunnel. The only information visible on the outer tunnel is the realm, as this is required for the routing through the proxies.

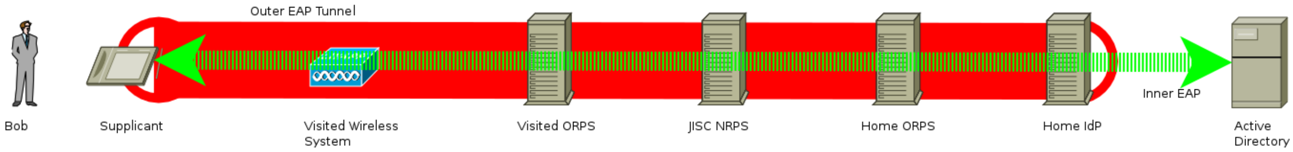

The end points of the inner tunnel are the supplicant and the IdP. The IdP needs to see the user's credentials so that it can perform the authentication with the data source. Thus this means that only the user and their home site ever get to see this sensitive information.

### EAP Types

Also, there are a lot of references to 'user credentials' rather than 'username and password'. This is intentional. The protocols which can be used with EAP are wide and varied. Because the inner tunnel is only visible at the two ends pretty much any form of credential can be used – username/password (EAP-PEAP), client certificate (EAP-TLS), SIM (EAP-SIM), token key (EAP-GTK)– whatever EAP supports. However, most sites use username and password (EAP-PEAP, or

PEAP for short), with a few using client certificates (EAP-TLS). As long as the outer ID contains the realm in the right format (“@holby.nhs.uk”) then the proxies will be able to route the request correctly.

### Topology

The design of this topology within a site depends on many factors. One of the most important is how your realms are structured. If a site has a single realm then when a request arrives at a site it'll have to be routed to an IdP. If there's a single AD server then that's easy. If the credentials are split across lots of AD servers then the IdP would have to try all of them. Thus it might be easier to have multiple realms with multiple IdPs which map to the AD servers more sensibly, e.g. one realm per AD domain.

Conversely it's entirely feasible that a site might have many realms but all the user credentials are contained within one AD server (or domain). In that case the ORPS has to proxy them all to a single IdP which is configured to handle all the realms.

### Resilience

It's a simplification to talk about 'a single IdP' or 'one ORPS' – in most scenarios sites would want to run multiple server or VMs for resilience. RADIUS requests can be load balanced between multiple servers either using internal load balancing or hardware load balancers (which can handle EAP statefully). The RADIUS servers themselves are effectively stateless and need no form of clustering. So 18 wireless controllers could generate requests to 4 ORPS which proxy to 6 IdPs or 3 top level proxy servers.

## Govroam Service

The expectation of the govroam service is that the site will use their existing wireless offering and configure up an SSID of 'govroam' using 802.1x. They will have, or set up, an ORPS through which authentication request are routed for govroam. An existing, or new, RADIUS server will act as the IdP to their AD/LDAP system.

The ORPS will be configured to communicate with Jisc's proxies both to send and receive authentication requests.

Once the authentication has happened the user's device will be put into suitable VLAN and assigned an IP address. The service specifies a (small) number of IP protocols and TCP/UDP ports which need to be open. The basic principle is that users of the service should have an expectation of a minimum level of network access on any govroam site. Generally this is web and VPN.

The RADIUS servers also need to log the authentication requests for a minimum period for audit purposes.

Sites should make an effort to ensure that their system is secure, monitored and maintained.

# User Authentication – More Detail

Ostensibly the Govroam service, provided by Jisc, is an authentication mechanism, no more, no less. However, on its own this is not a fully formed practical operation. In many cases it can be simply a layer that's added to an existing fully formed wireless service but in others there is a great deal of enabling work needed to make it work.

So what's required beyond RADIUS servers?

## Supplicant

The supplicant is the software that runs on the client device. Normally this is built in and integrated with the device e.g. on an iPhone it's the software that pops up the 'username/password' box or prompts you to accept the RADIUS certificate. The supplicant is responsible for managing the credentials and reacting to the prescence of an 802.1x connection. When a client connects to the Govroam SSID the supplicant initiates the 802.1x negotiation with the wireless system. Many supplicants can be configured using downloadable profiles.

## Wireless

Or, indeed, wired network switch. These devices are often know as Network Access Servers or NAS. The principles are the same in either case. A network switch and the wireless AP/controller are essentially both layer two devices that have a port a device connects to. The switch has a physical one, the wireless controller a virtual one. The NAS is configured to deny access to devices other than so that they can negotiate a 802.1x connection and, hopefully, an successful authentication.

The 802.1x communication between the client device/supplicant and the network equipment results in the NAS generating a RADIUS authentication request and is sent out to the next hop RADIUS server as configured above. If a positive reply is received the NAS opens the port, if negative leaves it closed awaiting another attempt. Disconnecting from the port cancels the authentication state for that device and a new connection will initiate another authentication attempt.

Both switch and wireless have to have their 802.1x profiles configured with which RADIUS server(s) to send the requests to. This is normally an ORPS on the local network. They need a network (VLAN(s), subnet(s)) onto which to drop the authenticated users. The wireless system needs an SSID of 'govroam' configured (802.1x/WPA2/AES). IP addresses need allocating (normally via DHCP).

## Network

Once the local NAS has opened the port to allow a device on the traffic is allowed to pass. The client appears on a VLAN (if a VLAN is the way it's configured), attempts a DHCP request (if that's the policy for the site), assigned an address (if it's appropriate for that device) and connects to the network. The size of the address pool should be about twice the size of the maximum number of concurrent users expected. It's preferred that public addresses are assigned to visitors but NATing is fine if this is not possible, which is likely to be most of the time.

## Authenticated Users

Here's a common scenario which will shed light on what it's like to have visitors on your network:

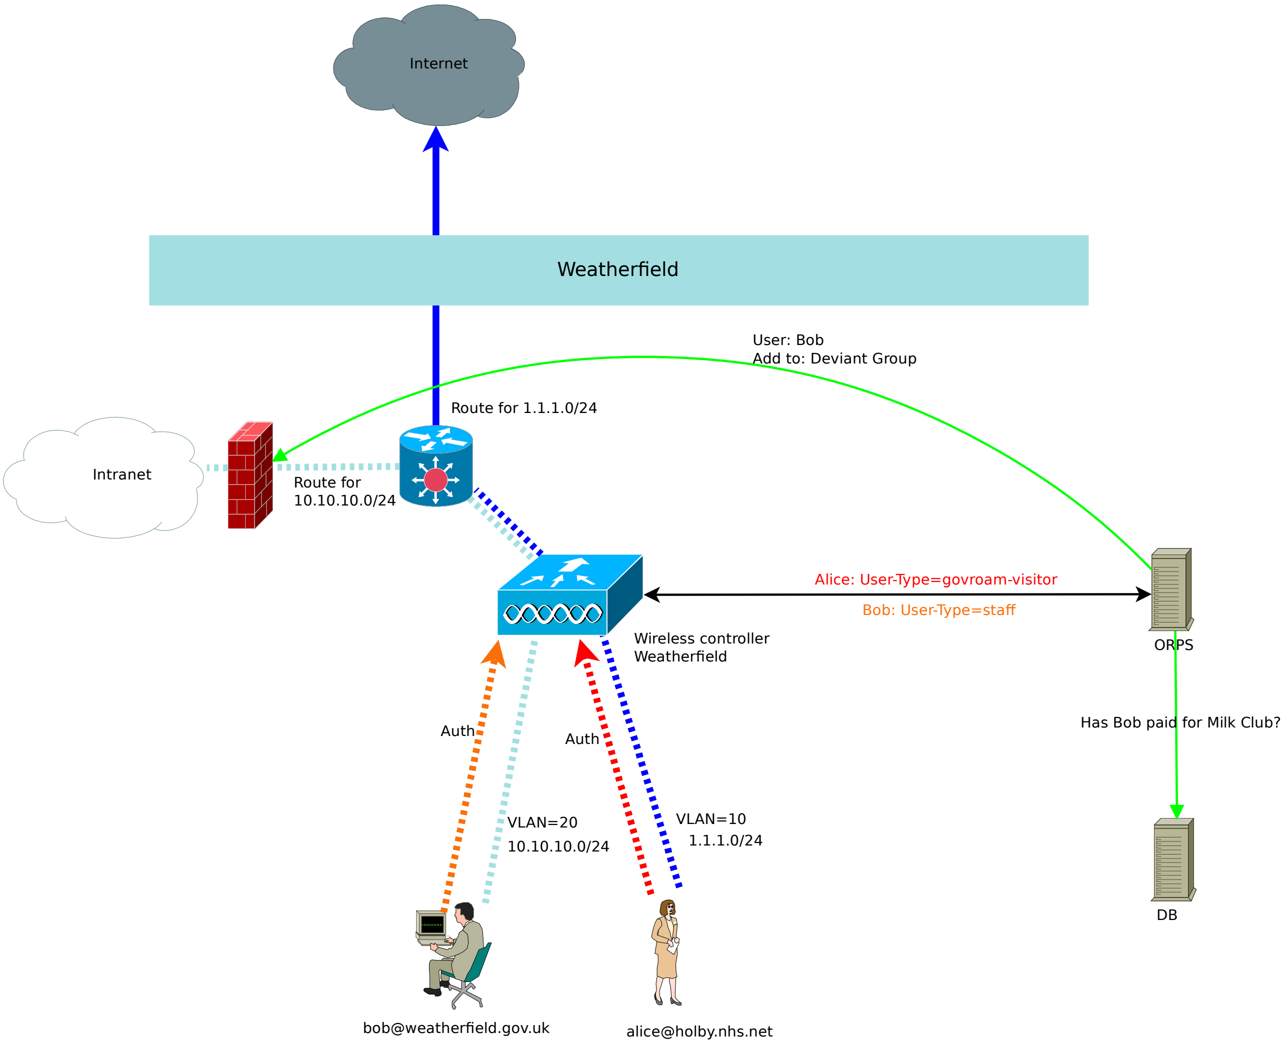

*Illustration 3: User Separation Scenario*

A Visitor from another site connects to your Govroam SSID and your RADIUS infrastructure can tell they're a visitor because the response came from the RADIUS proxies rather than your IdP. The RADIUS infrastructure sends back an attribute (which are normally proprietary) in the authentication response, which we'll call 'User-Type' for simplicities sake (you'll need to find which one is appropriate for your NAS). This attribute has a value of 'govroam-visitor'. Your NAS is configured to use User-Type to determine a VLAN (again the details of this tend to be proprietary). A value of 'govroam-visitor' might map to VLAN 10 whereas the default VLAN (if no User-Type attribute or if the attribute contains any value other than govroam-visitor) might be VLAN 20. The NAS applies this VLAN to the port that the client is on *before* it opens it. Thus when the client is allowed to connect it's already on its appropriate VLAN. So you now have a way of differentiating between users at layer 3 and thus can apply a different VLAN, different address range, different QoS, different firewall rules, different VRF, different internet connection etc.

In the eduroam community many Universities started by using eduroam in the most simplistic way – for visitors only. Then they realised that the above method meant that they could replace their local 802.1x SSID with eduroam, separate the users using RADIUS attributes, simplify the configuration, making roaming the default and reduce their SSID count, reducing their beacons and increasing the RF throughput. i.e. everyone connects to 'eduroam' and RADIUS does the work of figuring out who goes where.

Expanding on the above use of the RADIUS attribute, you can do anything you want as long as you have a way of differentiating users at RADIUS authentication time. Imagine a building that houses a number of different local services (NHS, fire, police, local government). If you configure the IdPs for each of these to send back a User-Type set appropriately (govroam-nhs, govroam-fire etc.) then you could setup each with their own VLAN/subnet quite easily just by configuring the wireless system as so.

Taking it further, then the better RADIUS servers have the functionality to query external data sources (databases being the obvious one but files, LDAP and AD are common) but also to run arbitrary commands at each point in the authentication process. e.g. post authentication, but before the response is sent, run a command that interacts with the API on your firewall sending the username of the person just authenticated with an instruction to apply a quarantine policy because of a flag set in a database.

## Policy

One of the underlying principles of Govroam is that visitors should expect the same network experience where-ever they are. Thus there is a set of ports which are expected to be open to the internet. Essentially these are VPN, web and email submission. Bandwidth shaping is allowed and fairly common. No one expects to have unlimited bandwidth but it ought be more than a trickle as people are expecting to be able to work from your site. Your local site policy may expect stricter controls but remember that this is a way of sharing so please try to treat visitors to your site as you'd want other sites to treat you when you visit them.

## Logging

Govroam's T&Cs require sites to log authentication information for a number of months for audit purposes, which is a fairly standard approach. However, there are a couple of other aspects which are worth understanding. There's a RADIUS attribute 'Operator-Name' which should be set to be the name of the site (format '1<site>', e.g. '1holby.nhs.uk'). Normally this is set by the ORPS and makes it much easier for home sites to know where their users are authenticating from for debugging purposes. Trying to audit and debug users activity by asking the administrator of each RADIUS server in a chain in turn is a difficult and time consuming process.

Operator-Name is also the prerequisite for a more specific user audit mechanism: Chargeable User Identity (or CUI). It used to be used in the dialup days to do what it says, charge users, but these days it's a way of identifying users in a secure way. When Bob visits a site their own logs won't know the contents of the inner identity, just the outer one, which is normally anonymous. So if Bob does something naughty and needs investigation or reporting then how does that site do that? Back to tracking users through a chain of RADIUS servers. Even with Operator-Name specified that's still potentially hundreds of users. The CUI attribute is populated with an encrypted hash of the user's Operator-Name and User-Name, giving a unique, but obfuscated, identifier for each user/site combination. Thus the visited site can send this information back to the home site who can then decrypt it to find the identity of their user and take appropriate action.

## Choices

Most of choices belong almost entirely to the site:

- RADIUS server hardware
- RADIUS server software – Jisc can offer some information on the options but preferably something that supports Operator-Name, Status Server and syslog.
- RADIUS server locations – one site or many, as long as they're all configured to handle authentication for the same domains.
- Server count – ideally two or more of each component
- Wireless system – anything which supports 802.1x
- Realms – format must be NAI RFC4282 and unique i.e. @something.somewhere.com
- Authentication source – normally AD but could be anything
- Additional ports and protocols – depends on local security policy
- VLANs
- Address range – public address space is generally best but NAT does work
- Internet connection
- Separation of home and visitors – home users at home can be treated as just 802.1x users and are not subject to govroam rules and regulations

## Configuration

There are about a dozen choices of RADIUS server and the configuration will be specific to a site, taking into account other uses and requirements of the wireless and RADIUS system. Offering definitive configurations is next to impossible but Jisc can offer general advice and some online resources.

# Glossary

ORPS – Organisational RADIUS Proxy Server. This acts as a 'router' for authentication requests. Authentication requests for known realms are routed to local IdPs. Those for unknown realms are passed on to a higher level to be routed.

NRPS – National RADIUS Proxy Server. Special case of an ORPS which is the top of the tree in a country and proxies between sites. Authentication requests for unknown realms are dropped.

RRPS – Regional RPS. Special case of an ORPS which sits at the top of a set of organisations and proxies between them. The level above would be the NRPS.

IdP – Identity Provider. Normally a RADIUS server with access to a store of credentials (normally Active Directory). Provides the actual 'Yes/No' (or Accept/Reject in RADIUS's case) answer to the authentication request.

Realm – a unique domain. This is used to identify the home authentication point for a set of users.

EAP – Extensible Authentication Protocol. A framework for handling authentication protocols. A number of EAP protocols exist for different types of credentials – username/password, certificates, SIM cards etc. Most commonly used are PEAP/MSCHAPv2 and EAP-TLS for username/password and certificates respectively.