CONTENTS

1. INTRODUCTION

This document provides summary info and guidelines for successful deployment of AudioFetch onto a network. This is not a comprehensive guide, though it is intended to cover most situations and be a quick read for those familiar with networking technology. If needed, further detailed information can be found in the AudioFetch App Note "Networking Details."

2. BANDWIDTH LOAD

Each AudioFetch user receives a separate audio stream from the AudioFetch box. Each stream consumes about 90 Kbps (including packet overhead). A maximum load of 250 simultaneous streams will consume about 22.5 Mbps bandwidth, which is why the 100base-T Ethernet connection on the AudioFetch box is adequate. Other wired portions of the network are recommended to be gigabit.

3. NETWORK CONFIGURATION GENERAL CONSIDERATIONS

- AudioFetch system will generally operate out of the box when connected to a simple network with one subnet such as created by a consumer grade wireless router using default settings, where no Guest Network is used for the WiFi.
- Access Point(s) of sufficient quantity and quality to support the expected user/traffic load should be used, enterprise-grade equipment will support more simultaneous users
    - If consumer equipment used, it should be high-end consumer-grade
    - Expect each consumer-grade Access Point (or wireless router) to support 25 simultaneous users in 2.4 GHz WiFi band. Configure for 5 GHz operation only to support larger number of users.
- Good WiFi signal strength/quality should exist throughout the facility

  - Areas of weak signal can cause just one user to consume lots of the WiFi bandwidth
- WiFi must support N mode, consider disabling B mode support (possibly even G mode)
- Access Points should be dual-band to reduce client load in crowded 2.4 GHz band
  - Disable 2.4 GHz band and operate only at 5 GHz if supporting more than 25 simultaneous users with consumer-grade equipment
- Access Points configured to minimize interference from other WiFi networks/channels
  - Double-check with a simple Android scanning App such as WiFi Analyzer or inSSIDer
- If Access Points are configured to use a "Guest Network" for end user access, devices connected to the wired portion of the network are generally not discoverable or accessible. In this situation:
  - Almost always, DNS discovery must be used (see information later in this document regarding "DNS LOOKUP" method.
  - Security settings must allow traffic between the guest network and the AudioFetch box residing on wired network.
  - Sometimes it can be helpful to configure the AudioFetch box to be part of the Guest Network using a Virtual LAN.
- Disable any automatic channel hopping feature (sometimes found in advanced access points), this will cause substantial disruption in the audio streaming
  - Or set the channel hopping to only occur rarely or late at night when users not present
- Enable WMM in the Access Points
- Access Points do not interpret QoS settings in IP headers consistently. AudioFetch packets are marked as VOIP per Cisco recommendations, some other brand Access Points will interpret as VIDEO.
  - Try adjusting the "Force alternate WMM VOICE QoS" setting in the AudioFetch Doghouse configuration – this will cause these other Access Points to interpret/handle the packets as VOIP
- If possible, enable per-user bandwidth limiting so that one user can't hog the WiFi bandwidth (with a large file download for example) – limit to what's reasonable for expected/allowed use
- DHCP server (to the mobile clients) should be set to provide lease-times longer than the usual/expected longest user audio session, as re-acquiring a lease could disrupt the audio

4. DISCOVERY

AudioFetch App running on user mobile devices must discover the IP address of the AudioFetch box at start up. This occurs easily on a simple network where the WiFi clients and wired (Ethernet) devices all run on the same subnet without any isolation or traffic restrictions between wireless/wired portions of the network, as is the case with a simple consumer-grade wireless router:

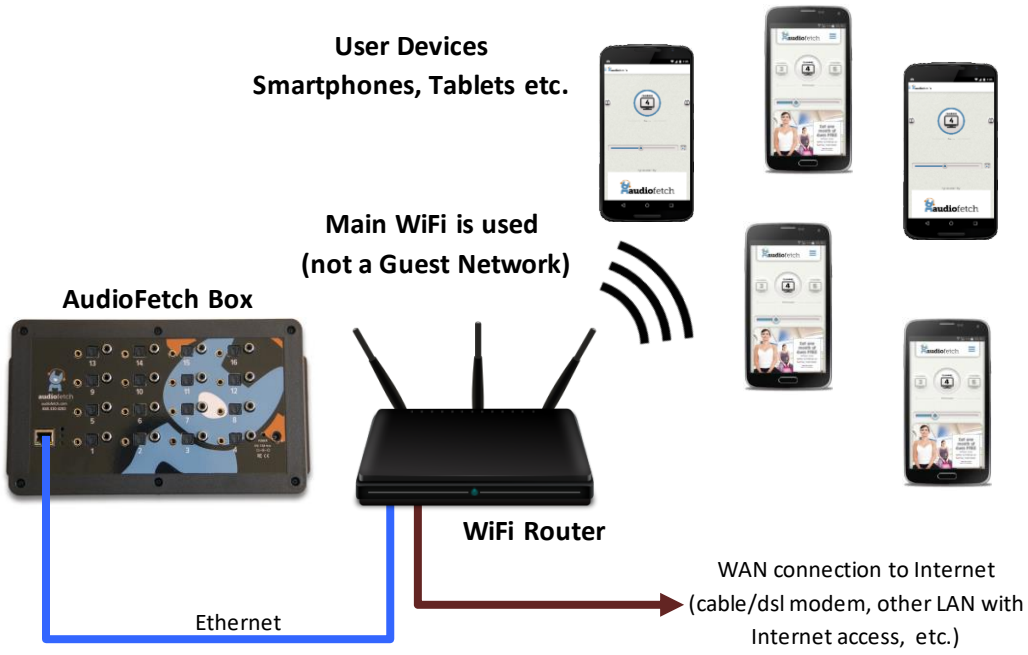

This network configuration is generally Plug-and-Play with AudioFetch.

If however the Guest WiFi network is used for customer/user access (most modern WiFi Routers support a main and a guest WiFi ssid), then the AudioFetch Box will be on the router's Primary network with all the users on the separate/isolated Guest network. The AudioFetch discovery traffic will be blocked by the security mechanism which isolates the Guest from Primary networks. With consumer-grade WiFi routers there will usually be no way to modify/configure security settings and in this situation one must usually employ a separate WiFi router for the customer/user WiFi (using its primary WiFi ssid) with the AudioFetch Box connected to it, and then the WAN port connected into the main network for user internet access.

Networks built on commercial-grade equipment with isolated end-user WiFi usually offer several different means to solve the discovery problem between end-user devices and the AudioFetch box. Such a network might look like this:

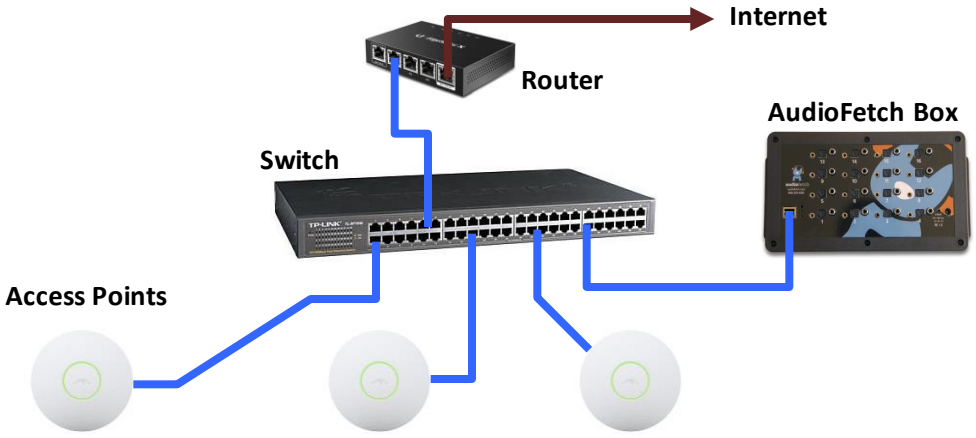

The Access Points might provide several different WiFi networks (customer, internal corporate, etc.) and often there will be security mechanisms implemented which isolate the customer WiFi from all other portions of the network and prevent customer/user devices from discovering the AudioFetch Box. There are several different solutions available in this situation:

- If Virtual LAN technology is supported, the easiest solution is to map the AudioFetch Box's Ethernet port into the same VLAN as the customer/user network.

- Alternately, if the Router supports local DNS Hosts entries, then:
  - The AudioFetch Box can be configured to a fixed IP address
  - A local DNS entry for the AudioFetch Box can be added to the Router
  - The Router must allow or be configured to route traffic between the AudioFetch box and customer/user devices on the WiFi network
  - See details for "DNS LOOKUP" in the next section

5. NETWORK TRAFFIC DETAILS - DISCOVERY

The AudioFetch App automatically attempts each of the discovery methods listed here, in the order listed, until one succeeds. Therefore a network need only support one. All traffic listed for a particular discovery method must be allowed on the network in order for that discovery method to work.

Discovery Method: SSDP

SSDP is a standard discovery mechanism employed by many types of equipment. Network traffic requirements for SSDP are:

Traffic direction: Mobile Device (AudioFetch app) → AudioFetch APB device:

Protocol	source port	dest port
UDP multicast (239.255.255.250)	p_1 *	1900
TCP unicast	p_2 *	80

Traffic direction: AudioFetch APB device → Mobile Device (AudioFetch app):

Protocol	source port	dest port
UDP unicast	1900	p_1 *
TCP unicast	80	p_2 *

* p_1 and p_2 indicate port numbers which have been chosen by the mobile device's operating system and typically are different each time the app runs

Comments on SSDP: This mechanism requires the network to allow multicast traffic from WiFi connected devices onto the wired portion of the network (where the AudioFetch APB is connected). In a complex network this traffic may also have to span subnets. Both of these are often challenges for larger networks where multicast traffic can be restricted for security reasons, and/or it is difficult to configure routers to pass multicast traffic between subnets.

Note about SSDP discovery on SECONDARY (management) port on AudioFetch Signature units with dual Ethernet ports: On these units the upper Ethernet port is called SECONDARY and is used to access the Doghouse configuration pages, which are not available on the lower/PRIMARY Ethernet port (used only for audio streaming). Purpose is to isolate and prevent access to the Doghouse by users. Discovery on the PRIMARY port works as described above. Discovery on the SECONDARY port works as described above except the UDP port number is 1901 instead of 1900.

Discovery Method: SSDP FALLBACK

SSDP FALLBACK is a modified version of SSDP where multicast packets from the WiFi-connected mobile devices are not required. Instead, the AudioFetch APB device (connected to the wired portion of the network) transmits unsolicited packets to a multicast address once every second. The idea here is that network security often allows multicast traffic from the wired portion onto the WiFi portion, but not vice-versa. After the mobile device receives the initial multicast packet from APB device, the balance of the discovery follows SSDP:

Traffic direction: Mobile Device (AudioFetch app) → AudioFetch APB device:

Protocol	source port	dest port
TCP unicast	p_2	80

Traffic direction: **AudioFetch APB device → Mobile Device (AudioFetch app):**

Protocol	source port	dest port
UDP multicast (239.255.255.250)	1900	1900
TCP unicast	80	p_2

Note about SSDP FALLBACK discovery on SECONDARY (management) port on AudioFetch Signature units with dual Ethernet ports: On these units the upper Ethernet port is called SECONDARY and is used to access the Doghouse configuration pages, which are not available on the lower/PRIMARY Ethernet port (used only for audio streaming). Purpose is to isolate and prevent access to the Doghouse by users. Discovery on the PRIMARY port works as described above. Discovery on the SECONDARY port works as described above except the UDP port number is 1901 instead of 1900.

Discovery Method: BROADCAST FALLBACK

BROADCAST FALLBACK employs broadcast packets from the WiFi-connected mobile devices in an attempt to circumvent situations where multicast doesn't work, somewhat of a "last ditch" kind of effort by the app:

Traffic direction: **Mobile Device (AudioFetch app) → AudioFetch APB device:**

Protocol	source port	dest port
UDP broadcast (255.255.255.255)	30981	30981
TCP unicast	p_2	80

Traffic direction: **AudioFetch APB device → Mobile Device (AudioFetch app):**

Protocol	source port	dest port
TCP unicast	80	p_2
UDP unicast	30981	30981

Note about BROADCAST FALLBACK discovery on SECONDARY (management) port on AudioFetch Signature units with dual Ethernet ports: On these units the upper Ethernet port is called SECONDARY and is used to access the Doghouse configuration pages, which are not available on the lower/PRIMARY Ethernet port (used only for audio streaming). Purpose is to isolate and prevent access to the Doghouse by users. Discovery on the PRIMARY port works as described above. Discovery on the SECONDARY port works as described above except the UDP port number is 30982 instead of 30981.

Discovery Method: mDNS one-shot

mDNS is a standard discovery mechanism employed by many types of equipment. mDNS "one-shot" is a subset of the mDNS protocol described in RFC 6762 section 5.1. Network traffic requirements for mDNS one-shot discovery are:

Traffic direction: **Mobile Device (AudioFetch app) → AudioFetch APB device:**

Protocol	source port	dest port	
UDP multicast (224.0.0.251)	p_1 *	5353	(source port must not be 5353)
TCP unicast	p_2 *	80	

Traffic direction: AudioFetch APB device → Mobile Device (AudioFetch app):

Protocol	source port	dest port
UDP unicast	5353	p_1 *
TCP unicast	80	p_2 *

* p_1 and p_2 indicate port numbers which have been chosen by the mobile device's operating system and typically are different each time the app runs

Discovery Method: DNS LOOKUP

Some networks simply do not allow any multicast or broadcast traffic in either direction between WiFi connected mobile device and the AudioFetch hardware device(s). This is especially true in situations where the two may be connected to different subnets, making it difficult to route multicast traffic between the subnets. Therefore, the DNS LOOKUP method may be employed which does not require any multicast or broadcast traffic.

It does however require manual entry(s) be added into the network's local DNS hosts list (usually maintained in the network's Router). It is assumed the reader is familiar with how to make these entries.

How DNS LOOKUP works in an AudioFetch system:

1. AudioFetch hardware device is assigned a static IP address on the network.
2. AudioFetch app running in mobile device sends a DNS query for the local host:
   "audiofetch.localdomain.localextension"
   A DNS lookup does not require any multicast/broadcast traffic and is almost always supported even on highly secured networks.
3. "localdomain.localextension" would be the connection-specific DNS Search Suffix provided to the mobile device by the DHCP server for the WiFi network. Easiest way to find this is connect a PC to the WiFi network and then run the command "ipconfig" from a command prompt, it should provide the "Connection-specific DNS Suffix" for the Wireless LAN adapter.
4. A local DNS hosts lookup table must be supported within the network (usually in the Router) and configured with an entry for this hostname (audiofetch.localdomain.localextension), which specifies the exact IP address of the AudioFetch APB connected to the network, the IP address does not need to be within the same subnet as the WiFi connected mobile device, however the network needs to be configured to route/allow traffic between the mobile devices and this IP address.
5. Note that some DHCP servers are not configured to provide a DNS Search Suffix, in these cases the DNS query sent by the mobile device will simply be:
   "audiofetch"
   in which case the local DNS hosts lookup entry would need to be: audiofetch Best case is to configure the local DNS hosts lookup to return responses for either type of query:
   "audiofetch.localdomain.localextension"

"audiofetch"

6. Mobile device uses the IP address returned in the DNS response to complete the discovery process.
7. Where multiple AudioFetch boxes are deployed in an installation, a larger set of entries in the local DNS hosts lookup table must be used. For a 2-box installation, the following entries are required:
   "audiofetch" (router should respond to either audiofetch or audiofetch.localdomain.localextension)
   "audiofetch-2" "
   For a 3-box installation:
   "audiofetch" (router should respond to either audiofetch or audiofetch.localdomain.localextension)
   "audiofetch-2" "
   "audiofetch-3" "
   For a 4-box installation:
   "audiofetch" (router should respond to either audiofetch or audiofetch.localdomain.localextension)
   "audiofetch-2" "
   "audiofetch-3" "
   "audiofetch-4" "
   (AudioFetch supports maximum of 4 DNS hostnames/entries)
   Note that there is no particular requirement to match a specific DNS name with a specific AudioFetch box/IP-address on the network, the only key requirement is that there be a unique entry for each and every AudioFetch box.
8. Note of clarification: There is no "audiofetch-1" entry used. The first box/entry is just "audiofetch" without the numerical suffix. Subsequent entries do have the numerical suffix such as "audiofetch-2", audiofetch-3, etc.
9. A maximum of 4 hostnames is supported by AudioFetch DNS Lookup discovery. For an installation with more than 4 AudioFetch boxes (up to 32 supported), additional boxes can be added to the "Discovery List" which is contained in the Doghouse configuration settings for each box. Further details may be found in the "Discovery List" section of the User Manual for your AudioFetch hardware.

Here is what needs to be done to configure your network correctly:

1. Add a static IP address reservation in your DHCP server settings for each AudioFetch hardware device connected to the network.
   a. NOTE: when the AudioFetch hardware device(s) request IP addresses from the DHCP server, they will advertise their own host names. THESE SELF-REPORTED HOST NAMES ARE NOT INVOLVED WITH THE DNS DISCOVERY PROCESS, PLEASE IGNORE THEM AND USE THE HOST NAMES DESCRIBED IN ITEMS 5-7 just above when configuring your router's Local DNS lookup settings.
2. In your network router's Local DNS lookup table, add the host names described in items 5-7 just above (add the appropriate number of host names depending on how many AudioFetch hardware devices connected to your network), and add the associated static IP addresses for each.

Normally the AudioFetch hardware device(s) do not require any configuration to make DNS discovery work.

Traffic direction: Mobile Device (AudioFetch app) → AudioFetch APB device:

Protocol	**source port**	**dest port**
DNS query (won't describe here)	(as usual)	(as usual)
TCP unicast	p_2	80

Traffic direction: **AudioFetch APB device → Mobile Device (AudioFetch app):**

Protocol	source port	dest port
DNS response	(as usual)	(as usual)
TCP unicast	80	p_2

Discovery Method: **QR Code IP Address – User Scans "Dogbone QR Code"**

This discovery method involves the listener scanning a QR code located on site that has the AudioFetch hardware's IP address embedded into the QR code. The QR code with the embedded IP address can be created and printed off by visiting www.audiofetch.com/qrcode. The QR code is branded with a "Dogbone" graphic that matches the graphic in the AudioFetch listening App so the user can easily correlate the QR code scan process within the App to the Dogbong QR code graphics placed through-out the listening facility.

6. NETWORK TRAFFIC DETAILS - AUDIO STREAMING

Audio streaming occurs after successful completion of discovery and primarily requires one-way UDP-unicast network traffic between the AudioFetch APB device and the WiFi connected mobile devices. Occasional bi-directional traffic is required for channel selection and keep-alive. None of this traffic requires multicast or broadcast, all audio streaming traffic is unicast:

Traffic direction: **Mobile Device (AudioFetch app) → AudioFetch APB device:**

Protocol	source port	dest port	
TCP unicast	p_3	6971	Control & keep-alive

Traffic direction: **AudioFetch APB device → Mobile Device (AudioFetch app):**

Protocol	source port	dest port	
TCP unicast	6971	p_3	Control & keep-alive
UDP unicast	6970	6970	This is the audio stream

7. NETWORK TRAFFIC DETAILS - DOGHOUSE CONFIGURATION

Access to the AudioFetch "Doghouse" web pages (to configure APB operation) requires the following traffic:

Traffic direction: **Mobile Device (AudioFetch app) → AudioFetch APB device:**

Protocol	source port	dest port
TCP unicast	p_2	80

Traffic direction: **AudioFetch APB device → Mobile Device (AudioFetch app):**

Protocol	source port	dest port
TCP unicast	80	p_2

Note that discovery is not required for access to the Doghouse configuration web pages, if the IP address of the APB device is known then it may be entered directly into a web browser for access. However the provided AudioFetch Doghouse Discovery software does rely on the standard discovery methods outlined in this document (currently it does not support the DNS method).

8. NETWORK TRAFFIC DETAILS - ONLINE FIRMWARE UPDATES

Access to online firmware updates requires the following traffic:

Traffic direction: **AudioFetch APB device → Internet:**

Protocol	source port	dest port
DNS (UDP packets)	53	53
TCP unicast	63713	80
TCP unicast	63714	80

Traffic direction: **Internet → AudioFetch APB device:**

Protocol	source port	dest port
DNS (UDP packets)	53	53
TCP unicast	80	63713
TCP unicast	80	63714