APEX Specification 0119 - Specification for the Passenger Connectivity Experience - QoE and QoS metrics

Passenger Experience Subcommittee of the Connectivity Working Group

Contents

Subcommittee Members

Blaz Vavpetic		blaz.vav@gmail.com
Peter Sevcik	NetForecast	peter@netforecast.com
Mary Rogozinski		maryrogozinski@wowway.com
Bernd Knobe	P3-Group	Bernd.Knobe@p3-group.com
Christian Schneider	P3-Group	Christian.Schneider@p3-group.com
Markus Michelt	P3-Group	Markus.Michelt@p3-group.com
Justin Eichenlaub	P3-Group	Justin.Eichenlaub@p3-group.com
Alan Jones	NetForecast	alan.jones@netforecast.com
Rich Evans	NetForecast	rich.evans@netforecast.com
Rebecca Wetzel	NetForecast	rebecca.wetzel@netforecast.com
Chimbu Aravind	Alethea	chimbu@alethea.in
Andrew Popov	Keyos	andrew@keyos.com
Rich Salter	The Salter Group	rsalter23@gmail.com
Stephan Schulte	Lufthansa Group	stephan.schulte@dlh.de

1. Introduction

Connecting Portable Electronic Devices (PEDs) such as mobile phones, tablet computers, laptop computers, and MP3 players to in-cabin wireless networks to access both onboard and terrestrial-based services is now a business-critical requirement for airlines. Passengers no longer view it as a luxury. Survey results and social media reflect the importance of Inflight Connectivity (IFC), as well as passenger frustration when the experience fails to meet expectations.

Insight into the passenger experience cannot be extrapolated from infrastructure-driven service level agreement reports of high performance from IFC providers because the actual passenger experience depends on factors that are not routinely monitored. These factors range from the process as well as the interface a passenger uses to attach to the service, to how they become entitled to premium services (e.g., free Internet connectivity, higher data speeds, or larger data limits), to how the service performs immediately upon attachment and throughout the journey.

The objective of the Passenger Experience Subcommittee of the Connectivity Working Group (via this document) is to identify the primary factors impacting the passenger experience of Inflight Connectivity (IFC) for offboard communications and Inflight Entertainment (IFE) for onboard entertainment. This will help airlines: understand and evaluate approaches to measure these factors; analyze the results to predict customer satisfaction; and tune their services to deliver the best passenger experience.

In reading this document, it is important to distinguish between the actual passenger experience (referred to here as Quality of Experience (QoE)) and performance of the underlying IFC-enabling infrastructure (referred to here as Quality of Service (QoS)). This document describes QoE performance factors that influence passenger experience quality as they connect, and once they are connected to the on-board network, and provides guidance for correlating QoS performance measurements from infrastructure components along the network path to the QoE performance factors.

This document describes metrics to evaluate IFC QoE and infrastructure QoS. Many approaches can be used to measure QoE such as: probes, vetted volunteers, crowd sourcing, apps on PEDs, and packet sniffing of real user traffic. Many approaches can also be used to measure QoS technical properties. How to measure QoE and QoS is the purview of airlines and their vendors, and therefore is not covered in this document.

1.1. The Connected Passenger Experience

The Connectivity Working Group is focused on the connected passenger experience offered on board. This includes the quality of the connectivity experience while boarding, attaching to in-cabin Wi-Fi, navigating the portal, acquiring Internet service, and once connected. Passenger interactions with airlines relative to connectivity and communications do, however, extend beyond the on-wing experience—thus this section also introduces aspects of the pre-flight journey that can influence a passenger's on-board experience.

1.1.1. *Pre-Flight*

A passenger's perception of an airline's digital performance begins with booking, usually through a web browser or a smartphone application. The booking experience is influenced by service responsiveness, ease of navigation and—of course—booking success.

As departure approaches, pre-flight communications for such things as check-in, flight status, gate information, congestion alerts, and on-board services facilitate the passenger experience. On-board connectivity information and purchase options are often included in these communications.

An airline can actively promote on-board connectivity to increase passenger awareness and manage expectations for in-flight connectivity. High awareness before boarding helps increase service adoption and ensure a seamless passenger connectivity experience. A passenger who does not know in advance how to connect to the service is more likely to have a negative experience and not use the service. Therefore, airlines should understand the pre-flight passenger experience and use that information to craft effective communications strategies.

Appendix A outlines best practices for pre-flight communications about IFC services to increase passenger awareness and satisfaction. Communication strategies should be tailored to passenger demographics and characteristics.

Because service awareness influences passenger experience quality, airlines should monitor the effectiveness of pre-flight communications. A range of indicators can be monitored to assess and optimize communication strategies and increase passenger satisfaction.

Once at the airport, passengers expect easy and cost-effective Internet connectivity. Recognizing this, many airports now offer free Wi-Fi. Although the airlines' role in this aspect of the passenger experience is beyond the scope of this document, it is an important part of the total passenger experience.

1.1.2. *Transition (Boarding the Aircraft)*

Upon entering the cabin, the passenger connectivity context changes. While still likely attached to a terrestrial cellular network and possibly still connected to the airport Wi-Fi network, passengers routinely take their seats and check e-mail, social media, and/or text message before takeoff. Airport Wi-Fi coverage may appear to be available but no longer be usable inside the aircraft, and cellular service may be degraded due to the physical environment and concentrated demand for cellular connectivity. If the in-cabin Wi-Fi is advertised, devices which have 'remembered' the SSID may automatically attach to it, but if the inflight connectivity service is not yet available or users are not automatically entitled to Internet connectivity, the devices are in a proverbial Wi-Fi dead zone—connected to the in-cabin Wi-Fi network, but unable to reach the Internet. Although they may have access to inflight portal content, passengers' first onboard connectivity experience may be perceived as being disconnected.

A common expectation is that while the plane is at the gate, whatever connectivity a passenger had prior to boarding should continue. Airlines should examine this use case and consider options, such as gate-to-gate connectivity, for mitigating this disruption. These options may include both marketing messaging and operational processes.

Appendix B contains a list of communication channels that can be used during the transition phase to promote and increase IFC service awareness.

1.1.3. *On Board*

This section applies to the passenger experience during the process of establishing connectivity to the in-cabin wireless network to access inflight entertainment and/or the Internet. In assessing the passenger experience, it is important to note that a passenger's experience may vary due to the capabilities and limitations of a passenger's PED.

The in-cabin connectivity experience once in flight falls into four categories:

- PED attachment to the Wi-Fi network
- Performance of inflight portal services
- Internet service acquisition (the process for purchasing, authenticating to, or otherwise acquiring service)
- Performance of the Internet service once acquired

1.1.3.1. *Attachment to an In-Cabin Wireless Network*

Attaching a PED to an in-cabin Wi-Fi network may require the passenger to take steps (perhaps repeatedly) that vary by Operating System (e.g., IOS, Android, Windows, macOS). When Wi-Fi service is made available, the on-board wireless network infrastructure must be able to accommodate a spike in the IP address requests (via Dynamic Host Configuration Protocol (DHCP)) that enable devices to communicate with the inflight portal for access to services. If the network cannot satisfy the requests, the device displays some variant of a "failed to connect to network" message and passengers again perceive they are disconnected. Passengers may repeatedly try to connect to the network with the same outcome until the congestion clears—resulting in a poor passenger experience.

Measuring attachment success rates and successful DHCP request fulfillments can provide an indication of passenger experience during this phase. While DHCP success can be measured from the wireless network infrastructure, instrumenting applications is necessary to automate the process of device attachment and quantify connection success or failure.

1.1.3.2. *Inflight Portal Services*

Once a device attaches to the in-cabin wireless network, passengers often access the inflight web portal to obtain flight information, services, entertainment, and the option to purchase Internet connectivity. The responsiveness of the web service that delivers the portal content via web browser or mobile application is next in the passenger connectivity experience. If the services load quickly, the experience is positive, but if the service is unavailable, terminates with error, or the passenger must wait for services to load, or the experience is bad the service has failed to meet expectations.

Passengers must log in using credentials that can take many forms: e.g., e-mail address, social network, digital certificate—and increasingly, the SIM card in the passenger's smartphone (a.k.a. Hotspot 2.0 or Passpoint). The passenger must accept conditions such as the use of personal data per the data protection laws of each country. In some cases, simply accepting conditions is sufficient to access the Internet without requiring additional data.

Portal content load times are good onboard Wi-Fi service performance indicators. When services are accessed via a mobile application, capturing and reporting the service response time or testing for page responsiveness can provide a credible local passenger experience measurement. Most web service management tool sets include ‘page load time’’ analytics.

Another portal experience indicator is to compare the number of devices that successfully attached to the in-cabin wireless network with the number of devices that successfully reached the portal. An airline can also compare the number of devices that reached the portal with the number of Internet sessions activated during the flight. Large gaps between the number of devices that reached the portal and the number of Internet sessions could mean that Internet session offers should be reviewed and adapted to the passengers’ expectation in terms of pricing or other service attributes.

Automating the attachment process with the mobile app and/or automating portal content delivery through an app provides a more consistent, reliable, and resilient passenger experience than web-based delivery. When a mobile app is required for IFC access but is not installed before the flight, it will negatively affect a passenger’s experience.

A number of other factors can also affect the passenger’s connection experience. For example, The service provider may set limits on such things as: upload and download speed, session duration, and maximum data usage. Service providers may also provide control over the applications (or ports) that can be used on their network. An airline may choose to block video streaming and voice calls, Finally, inflight portals typically use cookies in users' browsers to automatically re-log in on subsequent connections.

A Captive Network Assistant (CNA) can improve the passenger experience by speeding and simplifying IFC portal access. A CNA automatically redirects a passenger to the portal, thus eliminating the need to open a browser and trigger the captive portal manually with a search for a specific HTTP URL address.

The portal can also be used to manage passenger connection quality expectations by communicating information about the IFC network technology used on board and its expected performance during the flight. An airline can also provide information about factors that can influence the passenger experience, i.e., the number of currently connected users—and an airline can provide recommendations for users to maximize their experience.

The portal is a useful communication channel for optional services such as destination booking, on-board food and beverage ordering, leisure, e-commerce, and advertising. Some passengers may choose to spend time on the portal without connecting to the Internet. Measuring time spent on different portal services and pages can reflect the passenger’s portal experience. If a portal offers e-commerce, the number of purchases is another portal experience quality indicator.

1.1.3.3. *Internet Service Acquisition*

To acquire Internet service, passengers typically navigate the inflight portal until they can select—and if required pay for—Internet service. Whether the Internet connectivity is free or not, the ease with which passengers can acquire Internet service is important to the overall connectivity experience. Airlines should consider all potential barriers for domestic and international travelers and make this process as simple as possible. A key to a satisfactory service acquisition experience is to provide easy and rapid payment and authentication options. Measuring the time from authentication on the portal to successful purchase can reflect the passenger experience with payment and authentication options.

For fee-based Internet access, service selection and payment processes are significant sources of friction. Multiple service options (e.g., 1-hour, 2-hour, 100MB, 200MB, full-flight, messaging only, or full-broadband) force passengers to think about their needs and choose—thus the options and their presentation should be carefully thought through by airlines. Once service is selected, the payment method adds to the friction. "Pay with Points", "Pay with Credit Card on File", or "Pay with Airline Subscription" options require the passenger to have a frequent traveler account and be logged into it. Do they remember their membership number and password? When frequent traveler program membership is used to acquire Internet access, passenger experience quality during service acquisition can be measured by login success rates.

Using mobile apps that remember the passenger's details to log them in to their frequent traveler account or 3rd party subscription helps remove friction from this process—even if the passenger is not automatically entitled to Internet service and must make a purchase decision from within the app. Airlines should consider enabling their app to streamline this process.

Paying with a credit card requires the typical passenger to retrieve their wallet and enter the details on a small screen in a cramped space where others can see their personal information. Passengers may experience other obstacles such as mistakes entering credit card information, process errors, or software errors, all of which can frustrate passengers and even lead to purchase abandonment.

Other payment methods have their own complexities. For example, paying with a voucher requires the passenger to have obtained one in advance, and paying with a 3rd party subscription (e.g., Gogo, iPass, or Boingo), requires the passenger to know their username and password.

Digital payment methods such as Apple Pay, Alipay, and PayPal can reduce the number of purchase steps—and some digital payment methods also satisfy stringent online payment regulations for strong user authentication such as the EU's Revised Payment Services Directive (PSD2).

1.1.3.4. *Internet Service*

The passenger Internet service experience encompasses two phases: performance immediately following initial service acquisition, and subsequent performance throughout the service period (partial flight, full flight, or pre-determined data consumption amount).A passenger's perception of Internet service is more likely influenced by service responsiveness than effective bandwidth. The perception of responsiveness is enhanced when a service provider quickly confirms successful service acquisition (e.g., via a browser-based landing page or in-app message). Measuring the time from the Internet session opening request to the start of the internet session can reflect the quality of the passenger experience once service is acquired. Many factors cause experience quality to vary throughout an Internet session, and to maximize passenger satisfaction service providers should monitor performance and notify passengers via alerts or status updates when problems occur. This is important whether passengers pay for or enjoy complementary service. Service availability and performance influence satisfaction for all users—but travelers paying for service will have higher service quality expectations.

Because bandwidth must be shared by all connected passengers, restricting low-priority traffic types such as automatic software updates and device backups can improve the collective passenger experience and prevent passengers with usage-based services from quickly burning through their data allotment. These automated processes often begin as soon a passenger connects to the in-cabin Wi-Fi.

Some countries require Internet service to be turned off when passing through their airspace, and some routes have Internet "dead zones". By alerting passengers in advance, airlines can mitigate the negative impact on the passenger experience.

1.2. In-Cabin Performance Measurement

It is important for service providers to measure the service performance parameters described in this document initially and throughout a flight. What to measure and strategies for measurement are covered in Sections 3 and 4 of this document. Wherever possible, the service level a passenger receives (e.g., messaging only versus full broadband service) should be captured to provide context to evaluate performance.

2. Quality of Experience (QoE) versus Quality of Service (QoS)

As mentioned in the Introduction, in reading this document it is important to differentiate between measuring the airline passenger experience (QoE) and reporting on performance of the underlying infrastructure enabling inflight services (QoS). Here are the differences:

- **Quality of Experience (QoE)** describes the end-to-end experience of individual passengers. As shown in Figure 2 for IFE and Figure 3 for IFC service, QoE is measured at the PED level and may have different contexts for different passengers. Individually it can be used to correlate data collected for a connectivity experience with anecdotal or "soft" feedback, such as a customer service interaction or passenger survey results. In aggregate QoE quantifies the inflight connectivity quality experienced by all passengers.
- **Quality of Service (QoS)** in contrast, describes the performance of the infrastructure delivering Inflight Entertainment (IFE) and Inflight Connectivity (IFC). QoS is measured for one or more systems but IS NOT measured end-to-end and cannot capture the actual passenger experience. Service Level Agreements (SLAs) rely on QoS reports for component systems.

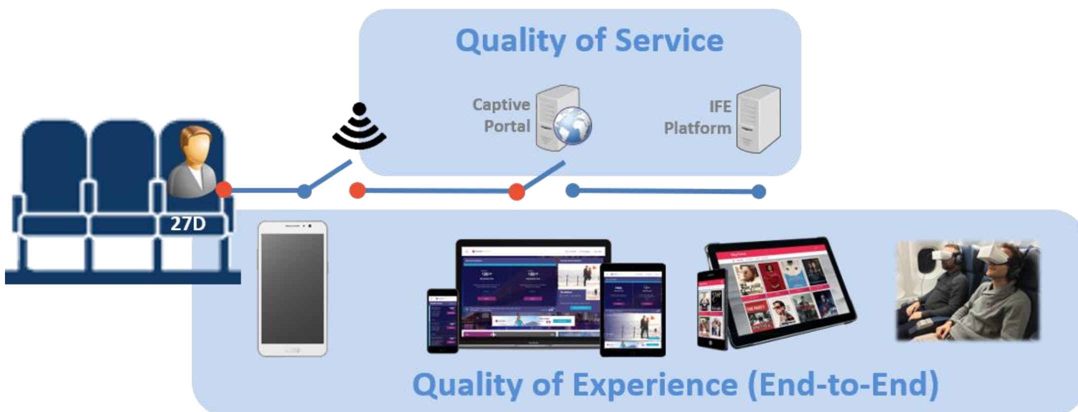

Figure 2. Inflight Entertainment End-to-End Connectivity

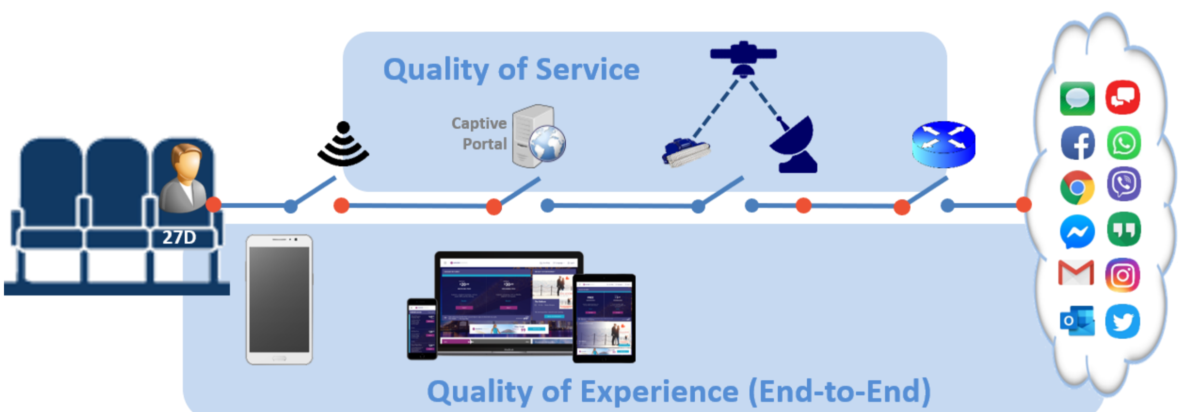

Figure 3. Internet Service End-to-End Connectivity

An example of the difference between QoE and QoS is the frequent discrepancy between service provider SLA reports (QoS data) that show good service uptime and availability, while passenger feedback (QoE data) characterizes service as poor. A passenger's experience is shaped by what they wish to do, their expectation of the service, and whether those expectations are met.

The typical applications passengers wish to use and sensitivity of those applications to changes in underlying technical characteristics are described in Section 2.

As a side note, Experience Level Agreements (ELAs) have been discussed as a more meaningful representation of delivered service quality than SLAs, however, ELAs require the ability to measure experience on a statistically relevant and diverse sample of passenger devices, which is difficult to accomplish.

What passengers say when rating or commenting on their *experience*

- "Pages took too long or did not completely load in my browser"
- "My Facebook feed wouldn't refresh—all I had was the spinning wheel at the bottom of the page"
- "I couldn't send an e-mail with a document attached"
- "I couldn't send or receive iMessage/WhatsApp/Facebook messages"
- "I kept getting disconnected from the service"
- "My device would not connect to the Wi-Fi"

2.1.1. *Expectation Management*

Managing passenger expectations when marketing IFC and keeping passengers apprised about service performance during a flight are important for passenger satisfaction. Passengers are somewhat tolerant of intermittent service interruptions if they are not prolonged, and the service provider provides feedback that it is aware the service is degraded or unavailable. The ability to alert passengers of an impending service interruption—as known from flight operations (e.g., satellite handoff) or historical data—is even more desirable.

Airlines may wish to enhance self-help tools such as FAQs on the inflight portal to help passengers recognize and to explain certain conditions. For example, an airline should explain that a passenger's experience may vary due to the capabilities and limitations of the passenger's PED, and/or they may improve passenger satisfaction by setting the expectation that experiences may vary due to network QoS policies (e.g., traffic shaping) that may affect some online activities such as streaming movies and live video. Also, airlines should clarify whether onboard Wi-Fi provides connectivity to the Internet as shown in Figure 3, or connects only to onboard entertainment as shown in Figure 2.

2.1.2. *Passenger Expectation by Passenger Type*

It may be possible to infer service expectations based on passenger type. Certain passengers may be more familiar with how service on an airline works and thus may not be as readily frustrated. Some may feel entitled to higher quality service based on cabin class. Others may be counting on service availability for professional productivity reasons. Some passenger classifications to consider may include:

- Frequent travelers versus occasional travelers
- Corporate travelers versus leisure travelers
- Business-class passengers versus economy-class passengers

2.1.3. *The Role of Service Marketing*

How service is marketed affects passenger expectations and perceptions. For example, if service is free, passengers are more likely to be happy with value received. Marketing a messaging service versus full high-speed Internet access, or promoting a heavily throttled service to ensure the ability to provide service to all devices may adversely impact passenger satisfaction. Airlines should design their service marketing based their passenger mix.

In any case, passengers may be dissatisfied, and airlines should monitor those experiences to understand and improve satisfaction levels.

2.1.4. *Other Factors*

Other factors influencing passenger satisfaction include:

- The ability to support multiple devices on a single purchased Internet plan
- The ability to continue using an Internet session purchased for a previous flight
- A consistent experience across an airline's fleet (i.e., a similar experience regardless of aircraft type or connectivity provider—common portal, payment methods, etc.)
- Availability of in-seat or seat-back power
- Differences in experience or performance between PED types
- Whether a passenger is a heavy or light user
- Comparisons with peer performance
- The influence of previous IFC experience
- Differences in expectations between "Digital Natives" (born during or after the onset of the digital era) and "Digital Immigrants" (born before the digital era).
- The applications used and the number of applications simultaneously active on a passenger's PED

Though an airline cannot account for all factors, they should strive to understand factors that may shape a passenger's service perception. For example, factors such as which seat a passenger occupies are important to correlate with service performance measurements.

3. Metrics to Quantify the Pre-Connection Passenger Experience

This section applies to the quality of the passenger experience when connecting to the Internet. It covers passenger experience quality metrics that can be measured and evaluated independently of the airline, aircraft, entertainment service, communications technology, and communications vendor. These metrics can be used to quantify and improve the passenger's pre-connection experience.

The pre-connection quality metrics are impacted by four systems in the wireless access network, as shown in Figure 4:

- **The Wi-Fi system**, with its authorization and association frames that enable a user to connect to the Internet.
- **The DHCP server**, responsible for giving an IP address to new requests from PEDs.
- **The DNS server,** which translates Web addresses (intelligible to users) to public IP addresses.
- **The RADIUS server**, responsible for authenticating user credentials in its database.

The ways in which the behavior of these four systems can degrade the user experience are varied and complex, and are therefore beyond the scope of this document.

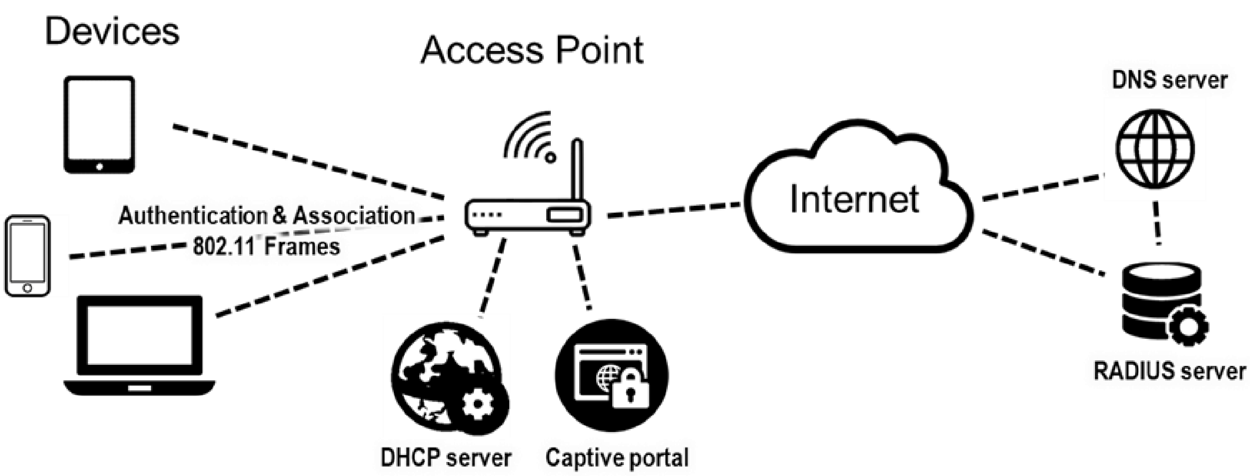

Figure 4. Infrastructure Supporting the Passenger Connection Experience

3.1. Network Metrics that Impact the Passenger Connection Experience

This section defines the metrics that affect the pre-connection quality of experience.

3.1.1. *Time to Associate (in seconds)*

Time to associate measures the number of seconds between the time the passenger selects the network to which they want to associate, and the passenger receives the association response from the Wi-Fi network.

3.1.2. *Time to Authenticate (in seconds)*

Time to authenticate measures the number of seconds from the time the user starts the authentication process until the process completes.

3.1.3. *Number of Steps (clicks/taps)*

The *number of steps* metric measures the number of clicks that the passenger must complete to navigate through the menus and forms on the captive portal to connect to the Internet.

3.1.4. *Time to Connect (in seconds or minutes):*

Time to connect measures the seconds or minutes that pass between the time the passenger selects the network with which they wish to associate, and when they successfully connect.

3.1.5. *Average Time to Abandon (in seconds or minutes):*

The *average time to abandon* measures the time that passes from when the passenger selects the network with which they wish associate until they abandon the effort.

3.1.6. *Automatic Browser Opening (yes/no)*

The *automatic browser opening* metric is a binary measurement that captures whether the web browser automatically requests the passenger’s credentials, or the passenger must open the browser.

3.1.7. *Successful Connects (percentage)*

The *successful connects* metric is the ratio of successful to attempted connections.

3.1.8. *Unsuccessful connects (percentage)*

The *unsuccessful connects* metric is the ratio of failed to attempted connections.

3.1.9. *Rejection Ratio*

The *rejection ratio* is the relation between the number of passengers who abandon their attempts and the total number who try to connect.

3.1.10. *Average Number of Retries*

The *average number of retries* metric captures the number of times the same passenger tries to reconnect to the Wi-Fi network until they succeed or abandon.

4. Metrics to Quantify the Passenger Experience Once Connected

This section applies to communications quality from passenger devices after the passenger has access to either IFE or IFC. Passenger experience quality can be measured by "soft" feedback obtained by polling or observing passengers, and by "hard" measurements of the actual passenger experience. For additional insight, this soft and hard data can be correlated with information such as the number of simultaneous users, passenger ticket price, on-time performance, etc. This section covers six hard passenger experience quality metrics that can be measured and evaluated independently of the airline, aircraft, entertainment service, communications technology, and communications vendor.

4.1. Application Categories

The following categories of applications make similar demands on the communications path between user and server. Not only do these application categories share behavioral characteristics, the users within each application category share generally consistent expectations.

Applications can be further grouped based on their sensitivity to network performance.

- Messaging: text messaging, Skype text, WhatsApp, Messenger (Facebook), Slack, Snapchat, etc. (generally all asynchronous text-based communications)
- Interactive: Web, shopping, email, Gmail, single-player games (casino, strategy), etc. (typically browser-based, html, JavaScript, or client-server)
- VPN: The process of establishing and maintaining connectivity to a corporate or consumer VPN access to services *[Note: this metric applies to the "tunnel", not the applications traversing the VPN]*
- Cloud Services: Microsoft Office 365, Google G-Suite, remote SAP, Google Drive, Dropbox, Flickr, etc. (aka virtual desk infrastructure (VDI))
- File Transfer: software updates, backup, photo uploads, etc.
- Streaming: Netflix, Hulu, YouTube TV, Amazon Prime videos, YouTube, Pandora, etc.
- Social: Facebook, Instagram, Twitter, Reddit, etc.
- Realtime: VoIP calls, two-way video, Facetime, multi-user voice or video sessions, etc.
- Gaming: multi-player real-time games like Fortnite, etc.

The inflight passenger experience is based on deviations from baseline performance as seen by the knowledgeable user within the IFE or IFC use case. Passengers will select which applications they use most often on a flight, and their interactions with those applications will define their performance framework.

Unsatisfactory performance will have different manifestations for each application category as described below

4.2. Unsatisfactory Performance

Below are descriptions of some typical symptoms of poor performance caused by degraded communications performance. This is not intended to be an exhaustive list.

- Messaging: slow message sent confirmation
- Interactive: slow page load time
- VPN: dropped sessions, slow session recovery
- Cloud Services: slow screen reaction time, mouse tracking, keyboard-screen reflection
- File Transfer: transfer pauses, slow file load/transfer times, file corruption
- Streaming: video pixilation or loading interruptions, slow replay, slow restart from pause
- Social: time to access the initial content load, pause in the scroll to more content, or video streaming issues such as above
- Realtime: garbled voice, frozen video, pixilation, poor movement representation
- Gaming: slow player movement, high lag, other gamers 'beat you out', hard to stay alive

4.3. Mapping Application Categories to IFE and IFC

Not all the application categories described here are currently available or practical on commercial flights. We have described the full set of possibilities as they exist on today's Internet. Table 1 maps the applications to their use in IFE and IFC. It also serves as a guide for which applications may be enabled in the future as communications technology improves.

Table 1. Application Categories for IFE and IFC

			Use Cases	
			Entertainment IFE	Communications IFC
Application Categories	Request-Reply	Messaging	NA	Current
		Interactive	Current	Current
		VPN	NA	Current
		Virtual Desk	NA	Current
		Gaming	Current	Questionable
	Continuous	File transfer	TBD	Current
		Streaming	Current	Current
		Social	NA	Current
		Realtime	Questionable	Questionable

Legend

Current: Currently supported by some airlines.

NA: Not applicable since these applications require interactions with people or services on the ground.

Questionable: Realtime is phone or video calling. Phone calls have been banned in some places by regulatory authority. In the IFE case, it is unlikely that airlines will encourage seat-to-seat calls since passengers can already text chat from seat to seat.

TBD: To be determined.

4.4. Service Availability

If the service is unavailable, none of the other metrics matter or can even be measured. Availability is the inverse of unavailability, defined as no response to application events. Unavailability is caused by a failure along the communication path. Such failures cause great user frustration. As addressed in **Section 1.2.1**

Expectation Management above, where service interruptions are predictable (as might be the case for international flights passing through Chinese airspace), the ability to alert passengers to an impending service interruption improves their experience.

An operating IFE or IFC session may stop functioning during a flight after the authentication process was successfully completed (as described in Section 1). This is manifested by long application response delays. The application might indicate lost service connectivity. The onboard Wi-Fi service may inform users that the service is no longer functioning and force passengers to return to the Wi-Fi access authentication process.

Abandonment Rates Due to Unavailability. As mentioned above, availability should be reported throughout a flight. Long periods of unavailability may cause high service abandonment rates when passengers who initially connect give up after sustained periods of unavailability or service degradation. Abandonment indicates the worst possible passenger experience. Service providers should therefore monitor abandonment rates to fully understand the passenger experience. Abandonment is best reported from the system and not from passenger devices.

4.5. Network Metrics that Impact Application Performance

This section defines the metrics that support the application categories described above. Each metric has a different effect on the user experience within each application category.

The primary user expectation is that an "application service" responds soon after the user makes some form of entry (clicked, typed, spoken, etc.). An application service is a service operating on a server that is responding over the Internet to an application operating on a user's device. Responsiveness is important to users, and when a user describes an experience as "slow" they are referring to the aggregate effect of the following metrics (listed here in order of importance): latency, loss, and DNS lookup time. Jitter and bandwidth are also key to the performance of many applications.

4.5.1. *Latency*

Latency is how long it takes a packet to travel from the device to the server and back to the device (not necessarily the same packet on the return). Every application service incorporates many round-trip events. A web page is "built" by the browser following a script of HTTP Gets that are then replied to with some content to be displayed. A typical web page requires dozens of such events.

For QoE measurement purposes, latency is defined as either one-way elapsed time for a packet to travel (aka transit) from point A to point B, or the two-way round-trip time (RTT) from A to B and back from B to A. One-way latency is often measured across a single device (router, communications link, etc.). Two-way RTT latency is mostly measured across networks using ping tests or derived from monitoring TCP connection ACK times. It is very difficult to reliably measure one-way transit time over a complex network. Latency is reported in milliseconds (ms).

4.5.2. *Loss*

Packets may be lost traveling from device to server or server to device in any segment of the end-to-end communication. In either case, loss will cause the application—either on the device or server—to trigger another packet exchange to finish the information transfer. Lost packets are very costly to overall response time. The application on the device must first wait for the first RTT, then it must wait for either end to realize that a loss occurred, and again for a second RTT. Beyond application performance, packet loss represents unproductive traffic on the wire and can aggravate poor bandwidth conditions,

For QoE measurement purposes, packet loss is defined as a packet that is lost during an RTT event or within a TCP flow of packets. Packets may be lost traveling from device to server or server to device in any segment of the end-to-end communications path without knowledge of where they were lost. However, devices that discard packets can report the discard along with location and the reason for the discard. Packet loss is reported as the ratio of lost packets relative to packets sent in percent (%).

4.5.3. *DNS Response*

All Internet Protocol (IP) communication depends on a domain name-to-IP address lookup operated by the DNS system. Every endpoint (client or server) must interact with a local or regional DNS server at the start of every new or unique TCP connection. DNS query-response time is reported in milliseconds.

Most Internet site connections begin with a query to a Domain Name Server (DNS). For example, the domain name apex.aero is located at an Internet Protocol (IP) address of 104.154.36.108. The user's browser or application must ask for a name-to-address resolution (like a phone number lookup) to the DNS system. Internet Service Providers (ISPs) operate DNS to their users. Airline communications services likely offer the service as well. However, some devices or applications prefer to use third-party DNS services (e.g., Google Public DNS). Regardless of the DNS used, the query and response require another RTT before the "latency" RTT described above occurs.

4.5.4. *Jitter*

Jitter is the variation in one-way network transit time. Although similar to latency, it is really a measure of transit predictability. Several applications have built-in mechanisms that effectively cope with latency. Those mechanisms, however, rely on consistent latency. Jitter can be measured one-way or round-trip. Jitter is reported as transit time variation in milliseconds.

4.5.5. *Bandwidth*

Bandwidth is the capacity of a circuit, device, or network to transfer a volume of data (bits or bytes) from point A to point B over a fixed time period. It is fundamentally different from all the other metrics that measure transactional events: single packet success or time. Bandwidth measures the data volume over time in bits per second. There are two sub-definitions of bandwidth that apply here, provisioned bandwidth and effective bandwidth.

> Provisioned Bandwidth: Ability to continuously accept data volume for transmission to destination. This is typically used to applies to devices such as communications circuits, routers, and switches. The transfer rate (e.g., 10 Mbps) may be the full theoretical capacity of a circuit or the allocated capacity.

> Effective Bandwidth: Data volume successfully delivered to destination. This is the measure of user or application payload that is successfully delivered to the destination device over the data protocols used to move the payload. It is what a single user sees as the rate of information appearing.

Most users are accustomed to downloading large files/payload and letting the process run in the background while they do other things (usually activities that are not so capacity sensitive). Some applications, such as adaptive bit-rate streaming, adapt content to a slower server-to-device path by reducing video resolution (thus reducing the payload). Also, basic properties of the Transmission Control Protocol (TCP) that most applications use, limit maximum usable bandwidth per TCP connection. That limit falls as the Bandwidth-Delay Product (BDP) increases. (Delay here equals latency). Therefore, very high bandwidth may be imperceptible to users on long-delay networks.

Higher bandwidth moves content faster once it begins to move. This is an important distinction that affects user perception. Once the user makes an entry on their device, they will perceive two events. The first is the feedback that the entry was received and the "next thing" has started. The elapsed time between entry and response is mostly governed by the response time metrics.

It is therefore important to note the difference between quick-to-start (response time metrics) and speed-to-finish (bandwidth). Users react differently to the two depending on what they are trying to accomplish.

Note: There is a separate value to bandwidth when supporting multiple simultaneous user sessions (aka connections). In this case, bandwidth is capacity to support more online passengers. The total bandwidth within an aircraft (Wi-Fi) or to/from an aircraft (e.g., satellite) must be shared by the number of passengers using the service. Higher total bandwidth shared among more devices (e.g., in a larger aircraft) is beneficial, but it does not improve a single user's application experience under most conditions. Why? Because we assume there are more connections so the effective available bandwidth to anyone remains roughly the same. Since this document identifies factors to quantify the individual passenger experience (QoE), measuring the experience across multiple users is a QoS metric to be addressed separately.

4.5.5.1. *Bandwidth Measurement*

Provisioned bandwidth cannot be effectively measured from the PED. Only effective bandwidth can be measured, and then only by monitoring traffic flows on the device adapter from a position in the network stack or by generating and measuring the performance of load tests, which in themselves generate additional network load and adversely affect overall network performance.

Effective bandwidth measurements must be evaluated in the context of the provisioned bandwidth available to the passenger under the subscription they hold. 235Kbps of measured effective bandwidth on the downlink and 180Kbps of measured effective bandwidth on the uplink do not represent good performance for a 5Mbps/2Mbps high-speed Internet service, but they represent excellent performance for a messaging service that allocates 250Kbps to both the uplink and the downlink channels.

Available bandwidth varies based on the location of the resource a device is attempting to communicate with. Airlines can optimize on-board connectivity and performance to ensure a high-quality on-board services experience. They can also optimize the uplink/downlink channels between the aircraft and the edge of their network on the ground. They can even work with 3rd parties to optimize private peering and optimize routes to well-defined networks or resources, but they do not control performance of the public Internet or the services passengers access. Airlines should consider which services are most popular (e.g., Facebook, YouTube, their own airline website, etc.) and do what they can to optimize access to these resources from their ground stations.

4.6. Mapping Metrics to Application Categories

It is not necessary to know the absolute best values for each network metric. If a popular application category is providing satisfactory performance, the network metrics are sufficient to support baseline performance and we can deduce that the passenger experience is satisfactory.

Application performance deteriorates when the network metrics degrade from the baseline performance. The table below maps the sensitivity of each application category to degradation in each of the network metrics described above.

Table 2. QoE Network Performance Matrix (Application Category Sensitivity to QoE Metrics)

		Performance Metrics Application Sensitivity to Baseline Value Change				
		Effect of Latency	Effect of Loss	Effect of Jitter	Effect of DNS	Effect of Bandwidth
Application Categories	Messaging	Neutral	Neutral	Neutral	Low	Low
	Interactive	High	High	Neutral	High	Low
	VPN	High	High	Neutral	High	Low
	Cloud Services	High	High	Neutral	Neutral	Low
	File transfer	Low	Low	Neutral	Neutral	High
	Streaming	Neutral	High	Neutral	Neutral	Low
	Social	Low	Low	Neutral	Low	Low
	Realtime	Low	High	High	Neutral	Low
	Gaming	High	High	Neutral	Neutral	High

Legend
High – metric degradation has a noticeable negative QoE impact
Low – metric degradation has a moderate negative QoE impact
Neutral – metric degradation has no or negligible QoE impact

When evaluating user experience quality, response time is generally more important than bandwidth. It matters less how fast the task completes if the time to start the task is long. To illustrate, consider which condition is more likely to trigger you to leave a restaurant: the wait for a table or the wait for the meal?

4.7. Evaluating Passenger Quality of Experience (QoE)

This section generally describes measuring and analyzing the metrics to evaluate the passenger QoE for IFE and IFC.

4.7.1. *Testing the Metrics*

True QoE measurement must be made via active tests. To properly reflect the passenger experience, the testing methodology must be applied with a statistically valid sampling of the passenger population across a diversity of devices (brands/models) that reflects the general passenger device mix and a sufficient quantity of measurement points (flights/seats). Passive data gathered using any method has value but is insufficient and should be correlated with active measurement data.

As referenced in **Section 1.1.4.2. In-Cabin Performance Measurement**, data samples should be taken throughout all stages of the flight and correlated with the stage/position of the aircraft at the time taken. This correlation can enable airlines and their IFC partners to provide predictive guidance to passengers regarding the IFC service based on historical data. An example of such predictive guidance may be the proactive alert to service interruption as the IFC system executes a satellite handoff or the plane is entering a geographic area with no service.

4.7.2. *The End-to-End Principle*

The Internet is built upon a foundational “end-to-end” principle. Before the Internet, communications providers operated every part of a communications path, e.g., phone to network to phone. Similarly, many (older) aircraft communications systems are built on a single service provider delivering every part of the communications path. For these, service providers take on the responsibility for overall service quality via SLAs because they own and operate every element in the path.

The Internet is different. The end-to-end principle means passenger’s device on the aircraft and application services on the ground are the two ends temporarily connected over a vast array of communications subsystems, devices, paths, and suppliers to provide a user experience. No one vendor supplies every part of that path, and the path and its components change constantly during each flight. Each supplier may meet its SLA, but the total path may not be delivering what the passenger expected.

The goal of measuring the user experience cannot be achieved by gathering data from parts of the user-to-server communications path. It requires end-to-end tests between devices and servers. The measurements need not cover all devices or all servers, rather, they need to comprise a representative sample that engages all elements of the typical device-server path. as well as all the applications (see Section 2.6 above).

4.7.3. *Instrumenting a Flight*

Airlines need to enable automated instrumentation of end-to-end testing. This can be accomplished by placing probes on aircraft in strategic locations. However, it is easier and more representative of the user experience to have a smartphone app for passengers to install, and let the app measure their in-seat experience while traveling. An excellent distribution mechanism for such an app would be the airline’s own smartphone app.

Regardless of how the test app is deployed (dedicated probes or smartphones), the app should perform the same functional tests to provide a uniform basis of comparison. The tests must operate on an end-to-end basis with known servers on the Internet. IFE measurements will require testing to the onboard entertainment server. Each test should be designed to have a low impact on the aircraft-to-ground communications system. Tests must consume a very small portion of available capacity. Also, the tests need to have minimal impact on the user's device, so they do not interfere with the device use.

4.7.4. *QoE Measurements within Subnetworks*

The primary network performance metrics defined above—latency, loss, jitter, DNS response, and bandwidth—can be measured within subnetworks along the end-to-end communications path as well as from passenger to ultimate destination. Subnetwork QoE information can be useful to assess the performance of major segments along the complete network path.

Subnetwork measurements can be implemented separately or cumulatively. For example, a separate subnetwork test is between probes at both ends of an air-to-ground network. A cumulative subnetwork example is measuring (A) from a probe on an aircraft to the air-ground gateway on the airplane (Wi-Fi) plus (B) an additional test to the far end of the air-to-ground subnetwork. These two tests show the performance of A and A+B. Proper analysis can derive the performance of B without the need for separate probes on the air-to-ground subnetwork.

4.7.5. *Converting Measurements to Quality Assessments*

Performance measurement vendors are free to design their own test implementations within the guidelines of this document. It is envisioned that most tests will be some variant of a standard Internet ping test. How the metric measurements are gathered, analyzed, and reported is left to the measurement vendors. Best practices will convert measurement data into clear QoE results that are good predictors of Net Promoter Scores (NPS) used by airline executives.to gauge customer satisfaction and loyalty.

5. Correlating IFC QoS Data to the Passenger Experience

This section provides guidance for correlating quality of service (QoS) measurements from infrastructure components along the end-to-end network path to the QoE network performance factors identified in Section 2. This correlation provides predictive and diagnostic information to enable infrastructure vendors, service providers, and their airline customers to optimize the passenger experience.

To effectively correlate QoS measurements to QoE metrics, it is important to establish a common language for QoS measurements. This enables performance of the same technical properties to be tracked and compared across the entire network path. This section describes common terminology for technical properties that affect passengers' experience quality.

5.1. The End-to-End IFC Experience

The quality of a passenger's IFC experience depends on the aggregated quality of service delivered by all subsystems and devices along the network path. Not only does each system element contribute to the ultimate passenger experience, handoffs between system elements also influence the outcome.

The cumulative performance of all system elements along the data path—in the data center, through the internet, across an aircraft's satellite or cellular service, and through the onboard Wi-Fi network to the passenger's personal electronic device (PED)—influences the ultimate passenger experience (see Figure 4).

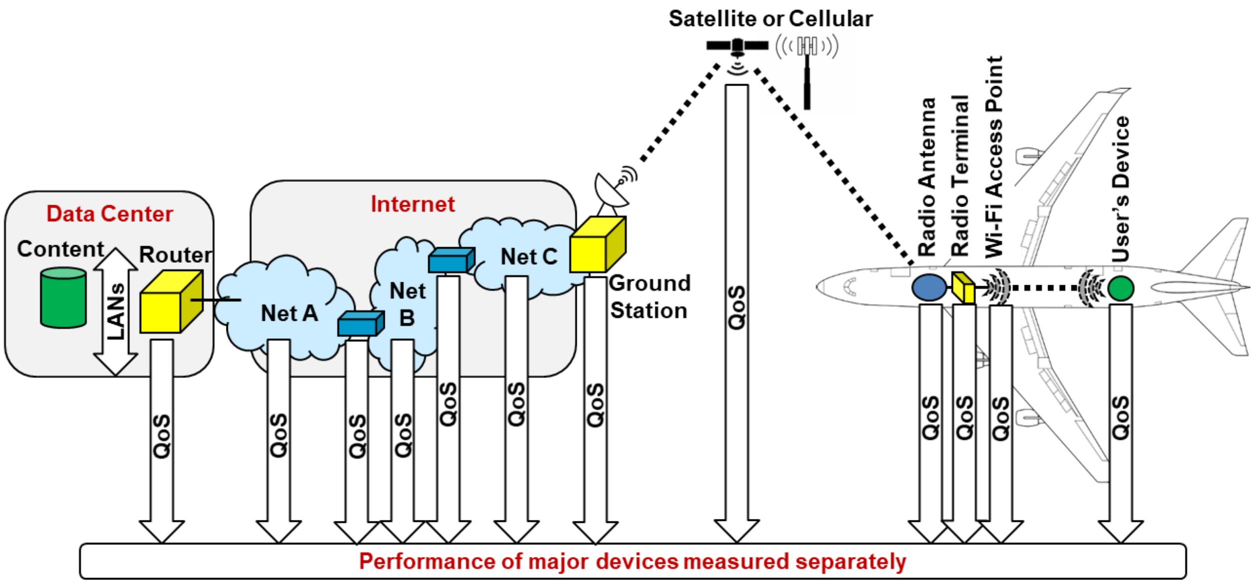

Figure 4. The QoS View of Network Performance

QoE is the end-to-end user experience between the user's PED and the "content" the user is interacting with somewhere on the ground—shown as the two green icons in Figure 4. The general approach for measuring QoE is described in Section 2 of this document.

5.2. QoS Measurements Along the IFC Path

To understand how each network component along the end-to-end path is contributing to the passenger experience, it is essential to correlate what is being measured along the path to the passenger's ultimate quality of experience. As described in Section 2.4, the key network metrics that apply to QoE are latency, packet loss, jitter, DNS response, and bandwidth. The same metrics also apply to QoS as measured at major components along the network path.

Nearly all components on an inflight communication path provide services with network metric properties that can be examined in this context. Proper documentation of just a few network metrics will show when QoE is poor based on a subset of degraded metrics and where QoS showed similar degradation to the same metrics.

It is important that measurements of these five metrics be as comparable as possible to correlate QoS information to QoE results. There are many sources of performance data and many ways to gather and analyze data feeds from various network devices along an end-to-end path. Detailed technical properties need to be assessed and integrated into these five general network metrics. In each case some annotation or qualification will be required (e.g., latency here is derived from the sum of transit delays for each communications device).

Telecommunications devices operate in many domains:

Analog: Fundamentally an analog signal is transmitted to a receiver.
Digital: The analog signal is interpreted into zero and one states known as bits then bytes.
Packet: A group of bytes (1 to ~1500) of payload with header and trailer. The header provides addressing and handling information for moving the payload from source to destination.

Each domain has its own version of the metrics in Section 2.5 plus many more that do not apply here. This document only applies to performance data in the packet domain. Although the same metrics in the analog and digital domains can impact performance in the packet domain, they are outside the scope of this document.

In some cases, infrastructure providers and network service providers directly measure the QoE network metrics described above—often, however, they measure other technical properties that contribute to QoE outcomes. Table 3 shows the network metrics' sensitivities to common communication device technical properties. Measuring these technical properties and assessing the results in an end-to-end context, makes it possible to correlate the measurements to the passenger's experience quality, and to diagnose network performance problems when they arise.

Correlating QoS measurements along the network path requires applying common terminology for those measurements across infrastructure and service providers. Table 3 presents suggested terminology for QoS technical properties that affect passenger experience quality. This terminology provides a starting point that is expected to evolve and expand over time in response to collaborative industry input.

Table 3. QoE Network Metric Sensitivity to Communication Device QoS Technical Properties

			QoE Network Performance Metric Sensitivity to QoS Technical Property				
QoS Technical Property	Units	Network Component	Latency	Loss	Jitter	DNS	Bandwidth
CPU Utilization	Percent	Device	Neutral	Neutral	Neutral	Neutral	Low
Memory Utilization	Percent	Device	Neutral	Neutral	Neutral	Neutral	Low
Active TCP Connections	Number	Device	Neutral	Neutral	Neutral	Neutral	Low
TCP Pool Utilization	Percent	Device	Neutral	Neutral	Neutral	Neutral	Low
Transit Delay	Seconds	Device	High	High	Low	Low	Neutral
Queue Depth	Bytes	Device	High	High	Neutral	Neutral	Low
Packet Discard	Percent	Device	High	High	High	Neutral	Neutral
Provisioned Bandwidth	Bits/Second	Link	Neutral	Neutral	Neutral	Neutral	High
Link Utilization	Percent	Link	High	High	High	Neutral	High
Link Latency	Seconds	Link	High	High	Low	Low	Neutral
Packet Loss	Percent	Link	High	High	High	Neutral	Low

Legend
High – noticeable QoS impact
Low – moderate QoS impact
Neutral – no or negligible QoS impact

5.2.1. *Communication Device Parameters*

Communication devices such as routers, gateways, switches, modems are single purpose computers with at least one processor chip (CPU) and associated memory chip(s). The processor and memory have a limited computational and data transfer capacity. The device vendor sizes these components to handle an expected set of real-time and background tasks under normal conditions. Real time refers to handling packets from an input port, making packet routing decisions, and handling the packet on an output port. Background tasks include updating route tables, system checks, and status communications with a management system. If device components are undersized (to reduce device cost, power consumption, heat, or device weight), the device may be unable to keep up with packet traffic load under typical conditions. Even if properly sized, the device may be unable to keep up under unusually heavy packet processing load.

If a device is operating near the limit of its traffic handling capacity it will present well known stress symptoms. A QoS monitoring system can detect these conditions if they are reported by the device on a regular basis. Of course, this adds more background processing to record these conditions and periodically transmit the readings to a monitoring system. These measurements are similar to those of the Task Manager utility in Microsoft Windows.

> **Anticipated Performance Effect of DNS over HTTPS**
>
> As of this writing, DNS over HTTPS (DoH) is emerging as a more secure alternative to conventional DNS. Its adoption will lengthen DNS lookup times because DoH performs DNS lookups using encrypted exchanges that require lookups to be performed at a few ground-based authorized DNS services. .

The following technical parameters that indicate potential impact to user traffic should be recorded and reported.

CPU Utilization – Current utilization of the main CPU, reported as percent utilization between 0% and 100%.

Memory Utilization – Current utilization reported as percent utilization between 0% and 100% of:

- Memory supporting the CPU
- Memory dedicated to storing user packets, reported as percent utilization between 0% and 100%

Additional performance parameters impacting user traffic are derived from the software operating within a communication device. The software likely has tables and policy on how packets are processed. Tables holding or referring to use packets (index to the actual packet stored in the device) may be essentially limitless (until running out of physical memory) or limited by software design or explicit policy.

Active TCP Connections – The current number of concurrent TCP connections or flows a device or server is handling reported as a number.

TCP Pool Utilization – If a device or server has a limit on the number of concurrent TCP connections or flows it can handle, the current utilization of the pool is reported as percent utilization between 0% and 100%.

Device Packet Transit Delay – Elapsed time between when a packet completely arrived on an input port and was placed on an output port queue. Essentially the time it took to determine the output port and place it there. This does not include time required for the packet to get to the head of the queue and be clocked out of the device. The transit delay is reported in milliseconds.

Output Port Queue Depth – The number of packets held in each output port. This is a clear indication of use traffic that is ready to move on towards a destination but “stuck” waiting for its turn to get out of the device. Queue depth is reported as the number of number of bytes in the queue. Some devices may not track bytes. If so, reporting the number of packets held in a queue is a less desirable but acceptable alternative.

Some devices are designed to operate with very large queues to prevent packet discard. This is prevalent in systems in which packet retransmission may adversely impact data delivery. However, such designs often present a phenomenon known as buffer bloat, which seriously degrades the performance of many applications. For this reason, tracking queue depth is an important QoS parameter in the QoE context.

Device Packet Discard – Devices have various algorithms that decide a packet cannot be processed or transmitted towards the destination. This is particularly a feature of firewalls. A discarded packet is a lost packet. QoS systems need to understand packet loss, which is typically attributed to loss on a communication link. Communication links do not report packet loss. The packet is simply gone. However, if a communication device discards a packet—for any reason—that is key data which provides insight into the cause packet loss. Packet discard is reported as the number of packets received but dropped by the device over the reporting period.

5.2.2. *Communication Link Parameters*

There is a link between every communication device pair. Some devices, e.g., Ethernet cables, are very reliable with predictable performance. But many in the context of IFC are complex with variable performance. The links themselves generally do not report performance (although there are some exceptions such as signal strength of radio channels). Typically, communication devices using links understand and report link performance. Therefore, despite having its own section here, link performance is an additional measurement and reporting responsibility of devices.

Provisioned Bandwidth – Most links have a fixed capacity to handle a set number of bits per second. This is also known as the clocking rate—the rate at which a 1,500 Byte packet can be placed onto a communication link. Once it is clocked onto the link and subsequently acknowledged by the receiving device, the sending device will remove the packet from its output queue. It is the theoretically fastest rate at which a packet can be handled.

If it is fixed, the parameter does not need to be updated in real time. However, if it varies over time or conditions (e.g., a satellite link changes as antenna orientation changes), it must be recorded over the reporting period. In all cases provisioned bandwidth is reported as bits per second in each direction (typically stated in million bits per second (Mbps)).

Link Utilization – A device knows the number of bits (or bytes) it clocked onto a link over a reporting period. This is a cumulative value based on how often a device performs the calculation. The link utilization is the ratio of bits sent divided by the theoretical bits that could have been sent over the same reporting period. Link utilization is calculated separately in each link direction and reported as a percentage. In simple terms, it is the percentage of provisioned bandwidth used over the reporting period.

Link Latency – In the provisioned bandwidth described above, there is a point at which the sending device finished clocking a packet onto a link and a subsequent time when it received an acknowledgement for the sent packet. The difference between those time stamps is the elapsed time for a packet to traverse the link to the receiver, the receiver to check that the packet arrived intact, transmit a short acknowledgement, and the acknowledgement arriving at the source. That elapsed time is the link round trip time (RTT). It is composed of two clocking times, one very short processing time (at the destination) and two propagation times. Generally, the propagation time is significantly longer than clocking and processing. The RTT/2 is therefore a good measure of one-way propagation time. Link latency is report as RTT in milliseconds.

Link Packet Loss – Many data link protocols have the ability to retransmit a packet that did not arrive at the receiving device (negative acknowledgement or no acknowledgement within a timeout period). If such a mechanism is operating over a link, the sending device knows the total number of packets sent and the number of packets needing retransmission over a reporting period. Link packet loss equals *(number of packets retransmitted)/(number of packets sent including the retransmissions)*. This is reported as a percentage over the reporting period.

6. How Network Performance Metrics Influence the Passenger Experience

As shown in Table 2, application categories have different sensitivities to degradation in each of the QoE network performance metrics. This means that a passenger's experience not only depends on network performance, but also on what they are doing. For example, streaming is very sensitive to the effect of packet loss, and under high loss conditions, a passenger's streaming experience is likely to be poor, while a passenger performing a file transfer may have an acceptable experience. On the other hand, during low bandwidth conditions a streaming user may have a good experience, while a passenger performing a file transfer may be frustrated.

Correlating QoS measurements to QoE outcomes involves examining the QoS data in a wholistic QoE context so their potential impact on the end-to-end passenger experience can be understood. Table 4 maps QoS technical property measurements to passengers' experiences using different application types.

QoE can be inferred from proper interpretation of network metrics. In the example in Table 4, .an increase in link utilization **(1)** causes elevated packet loss **(2)**, which can adversely affect the experience of passengers using interactive applications **(3)** because interactive applications are sensitive to loss. Messaging applications, on the other hand, are not as likely to be adversely affected by increased link utilization because they are generally insensitive to network packet loss.

The same process can be applied in reverse. For example, if file transfers are slow, it is useful to know that they are sensitive to diminished bandwidth, and the technical properties of system components that influence effective bandwidth should be investigated.

Table 4. Mapping QoS Technical Properties to QoE Outcomes

QoS

Technical Property	Latency	Loss	Jitter	DNS	Bandwidth
CPU utilization	High	Neutral	Neutral	Neutral	Low
BW provisioned	Neutral	High	Neutral	Neutral	Low
Link utilization (1)	High	High	Neutral	Neutral	Low
Queue depth	High	High	Neutral	Neutral	Low
Active Users	Neutral	Neutral	High	Neutral	High
TCP pool utilization	Neutral	Neutral	Neutral	Neutral	High

QoE

Application Class	Latency	Loss (2)	Jitter	DNS	Bandwidth
Messaging	Neutral	Neutral	Neutral	Low	Low
Interactive (3)	High	High	Neutral	High	Low
File transfer	Low	Low	Neutral	Neutral	High
Streaming	Neutral	High	Neutral	Neutral	Low
Social	Low	Low	Neutral	Low	Low

7. QoS Measurement Reporting Guidelines

The APEX Passenger Experience Subcommittee of the Connectivity Working Group recommends that QoS technical property reports encompass a minimum of 50 samples and include:

- Mean
- Median
- Minimum
- Maximum
- Standard deviation

This will govern the reporting period (per minute, hour, etc.). Continuous sampling over the course of a flight and the course of many flights will provide statistically significant information that will allow an airline to see the effects of such things as satellite or ground station handoffs, issues with onboard Wi-Fi equipment, etc.

8. Metadata Requirements to Support QoE/QoS Correlation

It is important that QoS reports include metadata that will enable correlation of QoE and QoS data. This metadata should include information about reporting intervals (e.g., every 5 minutes, 15 minutes, hour, etc.). It should also be time stamped and include location information (e.g., ground station, aircraft tail number, etc.).

9. Summary

An airline’s digital evolution and ancillary revenue success depends on keeping passengers connected both on- and off-wing. A passenger’s context and connectivity change constantly throughout each journey, and the more effectively an airline can remain in contact with a passenger, the better the passenger experience is likely to be. For example, timely information about such things as airport congestion, flight delays, gate changes, on-board services, airport maps, and baggage provided when and where needed personalizes and enhances the overall passenger experience.

Ironically, the highest risk of losing contact with passengers is in the place where the airlines have the most control—on board the aircraft. The perception of poor performance lowers the number of passengers likely to connect while on board, thus limiting an airline’s ability to remain connected to passengers while in flight, and thus missing opportunities to deliver additional services. To provide a pleasing experience and remove barriers to ancillary revenue, airlines require a well-conceived, reliable, and properly performing wireless service that offers a consistent fleet-wide experience.

Connectivity is now an important inflight service, therefore measuring and understanding the passenger experience is critical to predicting individual and aggregate passenger satisfaction. The measured experience can provide an airline and/or its service provider(s) with the data needed to identify service issues and address them as efficiently as possible.

Looking to the future, the communications infrastructure used to deliver IFC and IFE today will likely provide the foundation for additional services. For example, it may be possible to link passengers’ mobile devices to seat-back screens to access enhanced entertainment options, or to cabin crew systems that facilitate service personalization. Future services such as in-cabin IoT sensors may use the communication infrastructure to alert maintenance crews to proactively fix impending problems.

In conclusion, to improve and assure passenger satisfaction it is imperative that in addition to measuring the quality of service (QoS) delivered by IFC and IFE infrastructure, airlines and/or their service providers measure the quality of the **actual** passenger experience (QoE) at the PED level as described in this document.

10. Glossary of Terms, Acronyms & Abbreviations

APEX	Airline Passenger Experience Association.
Bandwidth	The theoretical capacity of a communication channel to transfer bits of data, typically measured in bits per second (bps).
Bandwidth Delay Product	In data communications, the product of a data link's capacity (in bits per second) and its round-trip delay time (in seconds).
Bit	A unit of electronic data.
bps	**Bits per second.** The rate at which bits move through a communication channel from Point A to Point B.
Byte	A unit of electronic data equivalent to 8 bits.
Communications Device	Router, switch, gateway, modem, radio, etc. along the network path between the passenger's PED and the server.
Communications Link	Data circuit, radio channel, inter-device interface, etc. along the network path between the passenger's PED and the server.
CWG	**Connectivity Working Group.** The APEX working group focused on enhancing the quality of inflight connectivity across the ecosystem.
DHCP	**Dynamic Host Control Protocol.** A network communications protocol which, among other things assigns IP address information to devices when they attach to the network.
DoH	**DNS over HTTPS.** A protocol for performing remote Domain Name System (DNS) resolution via the HTTPS protocol.
Effective Bandwidth	The observed or measured rate at which data is passed through a communications channel, typically measured in bits per second (bps).
ELA	**Experience Level Agreement.** Commitments from service provider to their customer on a measured set of user experience key performance indicators.
IFC	**Inflight Connectivity.** Services that connect airborne passengers to terrestrial-based communications networks (e.g., the Internet).
IFE	**Inflight Entertainment.** Services that connect passengers with entertainment available from servers on board an aircraft (e.g., movies, music, games, in-flight maps, etc...).

IFEC	**Inflight Entertainment & Connectivity.** A combined reference to IFC and IFE for aircraft that are equipped with or service providers to who deliver both.
KB	**Kilobyte.** Unit of data equivalent to 1,024 Bytes
Kbps	**Kilobits per second.** (meaning thousands of bits per second) is a measure of bandwidth and throughput (the amount of data that can flow in a given time) on a communications channel. Also expressed as Kbps.
MB	**Megabyte.** Unit of data equivalent to 1,024 Kilobytes (1,048,576 Bytes)
Mbps	**Megabits per second.** (meaning millions of bits per second) is a unit of measurement for bandwidth and throughput (the amount of data that can flow in a given time) on a communications channel. Also expressed as Mb/s.
NPS	**Net Promoter Score.** A customer satisfaction benchmark that measures how likely your customers are to recommend you to a friend.
Operating System	Also (OS). The software that runs on a Portable Electronic Device which dictates how the user interacts with the device for functions such as establishing connectivity to wireless network. Leading operating systems are Google Android, Apple iOS, Microsoft Windows, and Apple macOS, though others exist (e.g. Linux).
PaxEx	Passenger Experience
PED	**Portable Electronic Device.** Portable electronic equipment including but not limited to mobile/cell phones, electronic e-readers, tablet computers, laptops, MP3 players, and electronic toys.
QoE	**Quality of Experience.** A measure of the overall level of customer satisfaction with a service. QoE is related to but differs from Quality of Service (QoS), which embodies the notion that hardware and software characteristics can be measured, improved, and perhaps guaranteed.
QoS	**Quality of Service.** A network's ability to achieve its maximum performance and uptime through the management of bandwidth and other network performance elements like latency, packet loss.
SLA	**Service Level Agreement.** A commitment between a service provider and a client.
Speed	Perceived performance of a service (e.g. how fast a web page loads or how long it takes to complete a transaction).
SSID	**Service Set Identifier.** A sequence of characters that uniquely names a wireless local area network (WLAN).
Throughput	The rate of successful message delivery over a communication channel.

VPN	**Virtual Private Network.** A software defined secure connection between two hosts, typically a PED and a server, established over an unsecure communications channel such as the Internet.

11. Appendix A: Pre-flight Phase Communications Best Practices

As mentioned in section 1.1.1, this section outlines suggested practices to strengthen communication with passengers regarding IFC services. Each communication channel has a different impact on passenger awareness, and passenger receptivity to each channel will vary based on demographics.

	AWARNESS IMPACT		
Communication Channel	**Low**	**Medium**	**High**
Airline Mobile Application			✓
Airline Website		✓	
Email/Texts	✓		

Airline Mobile Application

Most passengers carry at least one portable electronic device (PED). Passengers with an airline mobile application on their PED can manage bookings, check flight status, receive upcoming departure notifications, and access a range of other information. Airlines are increasingly using mobile applications to upsell or provide additional services.

Airlines can promote IFC service on an upcoming flight by:

- using a mobile app to offer the purchase of Internet access before a flight. Once on board, the passenger can access the Internet using a promotional code or a dedicated login on the IFC portal (e.g., using seat number and last name).
- using push notifications (i.e., email and text messaging) to advertise IFC service availability. These notifications can be sent before boarding, while a passenger is connected to cellular or airport Wi-Fi service.

AIRLINE MOBILE APPLICATION	
Target Passenger Segment	Generation X & Y
Advantages	• Provides one channel to access Wi-Fi throughout the passenger's journey (e.g., airport/lounge Wi-Fi, in-flight/gate-to-gate Wi-Fi). • Facilitates a seamless and unified passenger experience. • Provides a one-stop shop for an airline to promote IFC services and provide information.
Disadvantages	• Requires passengers to install the airline's mobile application.

Airline Website

An airline can use its website to promote IFC service. Passengers are increasingly open to selecting add-on services during booking (e.g., special meals, seat upgrades, extra luggage, and Internet access). Many passengers are open to purchasing Internet access when they buy their tickets.

It is a good practice to state whether the aircraft has IFC service and to provide instructions on how to connect once onboard.

	AIRLINE WEBSITE
Target Passenger Segment	All
Advantages	• It provides opportunities to upsell or promote IFC services during passenger booking. • Coordinating between the mobile application and the airline website facilitates a more seamless passenger experience.
Disadvantages	• The airline website does not enable the airline to provide information shortly before takeoff.

Email/Texts

Email and texts can reach a large number of passengers across generations. Most passengers carry a PED, and this communication method is an easy and fast way to provide instructions and recommendations on available onboard service such as IFC (e.g., send a text with a hyperlink to the onboard IFC portal).

Looking ahead, Generation Z passengers will expect to communicate with airlines through messaging applications like WhatsApp and Facebook.

	EMAIL/TEXTS
Target Passenger Segment	All
Advantages	• Airlines can reach all passengers, as they generally have their phone numbers and email addresses. • These channels provide “last minute” information and can be used to communicate changes that may impact onboard services.
Disadvantages	• Texts may be perceived as intrusive and emails can easily be ignored by passengers.

12. Appendix B: Transition Phase Communication Best Practices

As mentioned in section 1.1.2, below are best practices to communicate with passengers about IFC services. Each communication channel has a different impact on passenger awareness. The effectiveness of each channel will vary based on passenger demographics.

	AWARENESS IMPACT		
Communication Channels	**Low**	**Medium**	**High**
Wi-Fi Onboard Stickers		✓	
Crew Announcement	✓		
IFE Video			✓
IFE Information	✓		
Airline Magazine	✓		
Airline Mobile Application			✓
Wireless IFE (W-IFE)		✓	

Wi-Fi Onboard Stickers

Informational seatback stickers showing a Wi-Fi icon are recommended to educate passengers about onboard IFC service, and to help them connect. A scannable QR Code can automatically direct passengers to the IFC onboard portal. A simple and easy-to-remember URL can also be used to access the IFC onboard portal.

WI-FI ONBOARD STICKERS	
Target Passenger Segment	All
Advantages	• Wi-Fi stickers inform all passengers about the availability of onboard IFC service.
Disadvantages	• None

Crew Announcements

Some airlines communicate about onboard services through crew announcements. This communication channel enables the airline to easily adapt messaging to circumstances. Crew announcements enable reactive and proactive communication about IFC service availability and quality during the flight, which can be helpful in managing passenger expectations.

	CREW ANNOUNCEMENTS
Target Passenger Segment	All
Advantages	• Enables messages to be customized to reflect current situations. • The announcements enable proactive and reactive communication to set expectations about service quality in real time.
Disadvantages	• Passengers may not pay attention.

IFE Video

An airline may choose to provide passengers in IFE-enabled aircraft with video instructions on how to connect to onboard Wi-Fi. As most passengers cannot access entertainment programming during announcements, an instructional IFE video is likely to capture passengers' attention.

	IFE VIDEO
Target Passenger Segment	All
Advantages	• Informs all passenger about IFC availability and use.
Disadvantages	• It may be expensive to produce and upload video.

IFE Information

Many passengers use their in-flight entertainment screens to view services and content. The IFE platform can be used to inform passengers about available IFC services, and provide access instructions.

	IFE Information
Target Passenger Segment	All
Advantages	• Provides a platform to inform passengers about available IFC services, and to instruct them about how to attach to Wi-Fi and access the Internet.
Disadvantages	• Not all passengers will view this information.

Airline Magazine

It is a recommended practice to devote a section of an airline's inflight magazine to informing passengers about available IFC services, and providing instructions on how to connect. An airline can also provide instructions and guidance to manage passenger expectations.

AIRLINE MAGAZINE	
Target Passenger Segment	All
Advantages	• Can be used to inform passengers who read the publication about the availability and use of IFC services.
Disadvantages	• Not all passengers read airline magazines.

Airline Mobile Application

An airline's mobile application can be designed to enable passengers to connect seamlessly to onboard Wi-Fi. Many passengers use airlines' mobile applications, and stored credentials can be used to automatically authenticate passengers and personalize the portal experience (e.g., customize the experience based on passenger status).

AIRLINE MOBILE APPLICATION	
Target Passenger Segment	All
Advantages	• Enables a seamless connection to onboard Wi-Fi. • Leverages the authentication mechanism available in the airline's mobile application. • Enables the airline to provide a personalized experience based on the passenger account details.
Disadvantages	• Requires integration between the airline's mobile application and the IFC service.

Wireless-IFE (W-IFE)

When aircraft are equipped with both W-IFE and IFC services, it is a good practice to promote the IFC service via the W-IFE service and enable passengers to use their PED to navigate from one service to the other. This can be facilitated with one-click access to the IFC portal.

	W-IFE
Target Passenger Segment	All
Advantages	• Can inform all passengers about IFC service availability.
Disadvantages	• Content will need to be updated.