NetForecast®

EVALUATING THE PASSENGER CONNECTIVTY EXPERIENCE

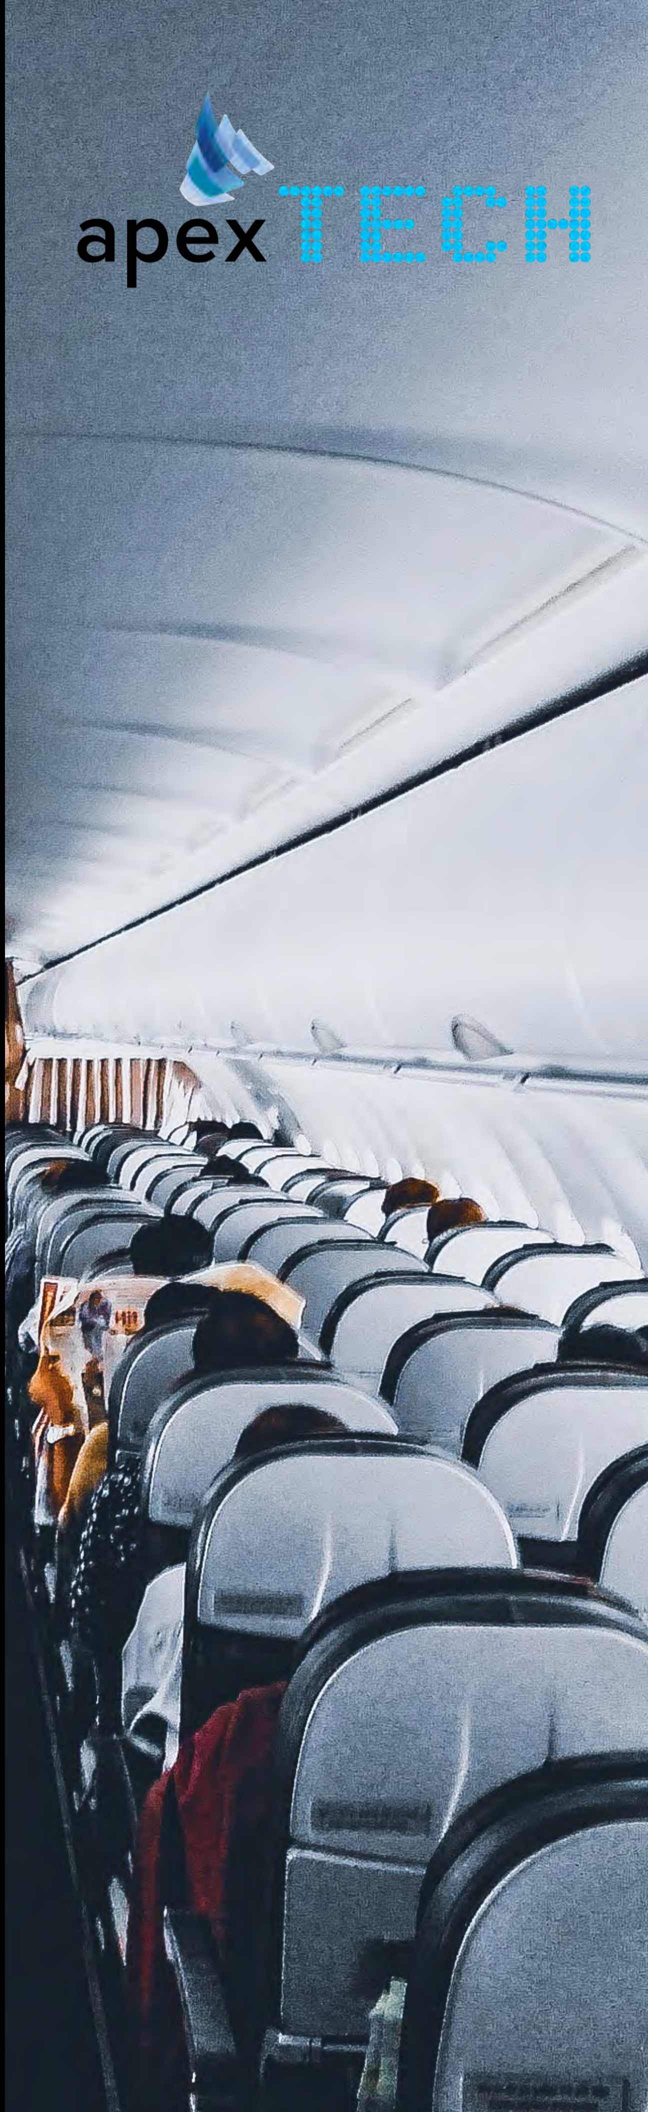

July 2019

Passenger Experience Subcommittee of the Connectivity Working Group

Evaluating the Passenger Connectivity Experience

Passenger Experience Subcommittee of the Connectivity Working Group

Contents

Subcommittee Members

Blaz Vavpetic		blaz.vav@gmail.com
Peter Sevcik	NetForecast	peter@netforecast.com
Mary Rgogozinski		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	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 a number of 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 document 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. This document describes end-to-end system performance factors that influence passenger experience quality (see **Section 2.7 Evaluating Passenger Quality of Experience**). How to measure these factors is the purview of airlines and their vendors, and therefore is not covered in this document.

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 addresses QoE. QoS is being addressed by the QoS Subcommittee of the Connectivity Working Group.

1.1. Connected Passenger "Journey"

Passenger interactions with airlines relative to connectivity and communications extend beyond the on-wing experience. Although the Connectivity Working Group is focused on the service offered on board, it is important to understand the onboard connectivity experience within the overall context shown in Figure 1.

Figure 1. The Total Passenger 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.

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

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 and 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 for mitigating this disruption. These options may include both marketing messaging and operational processes.

1.1.3. In Flight

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.

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

- Device attachment to the Wi-Fi network
- Performance of inflight portal services
- Process for purchasing, authenticating to or otherwise acquiring Internet service
- Performance of the Internet service once acquired

1.1.3.1. Attachment to an In-Cabin Wireless Network

Attaching a device to an in-cabin Wi-Fi network requires the passenger to take steps (perhaps repeatedly) that vary by Operating System (e.g., IOS, Android, Windows, macOS). When the plane reaches altitude and passengers are told "you are now free to use your personal electronic device in airplane mode", the on-board wireless network infrastructure must be able to accommodate the spike in 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. Performance of Inflight Portal Services

Once a device attaches to the in-cabin wireless network, passengers often access the in-flight 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 passenger must wait for services to load, the experience is bad, and the service has failed to meet expectations.

Portal and/or page load times are good indicators of service performance and the passenger experience. When services are accessed via a mobile app, capturing and reporting the service response time or testing for page responsiveness can provide a credible passenger experience measurement . 'Page load time' analytics are part of most web service management tool sets.

Like automating the attachment process with the mobile app, automating portal content delivery through an app provides a more consistent, reliable, and resilient passenger experience than web-based delivery.

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 to the passenger 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.

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? 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. Paying with a voucher requires the passenger to have obtained one in advance. Paying with a 3rd party subscription (e.g., Gogo, iPass, or Boingo) requires the passenger to know their username and password.

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.

1.1.3.4. Internet Service Performance

Excluding the in-cabin Wi-Fi network performance, 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).

1.1.4. Initial Service Performance

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).

1.1.4.1. Performance Throughout the Service Period

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.

1.1.4.2. 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 Section 2 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.

1.1.4.3. 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 seat 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.

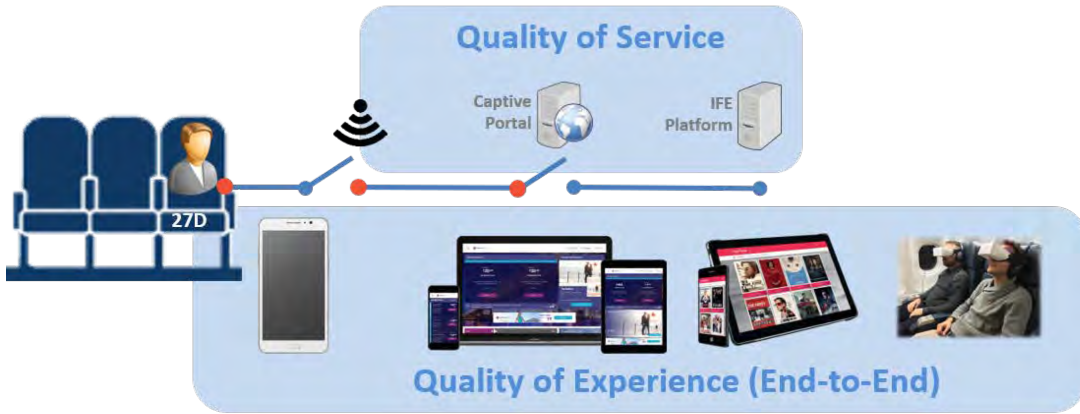

Figure 3. Inflight Entertainment End-to-End Connectivity

Figure 2. 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 99% 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 and their expectation of the service.

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"

1.1.5. Expectation Management

Managing passenger expectations when marketing IFC and keeping passengers apprised about service performance during a flight is 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.

1.1.5.1. 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

1.1.5.2. 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.

1.1.5.3. Other Factors

Other factors influencing passenger satisfaction include:

- The ability to support multiple devices on a single purchased Internet plan
- Availability of in-seat or seat-back power
- Differences in experience or performance between device types
- Comparisons with peer performance

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.

2. Metrics to Quantify the Onboard Passenger Experience

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.

2.1. Use Cases

To accurately assess the user experience, it is essential to have a framework within which to judge that experience. A user encountering an application for the first time may become frustrated due to poor operational knowledge rather than poor system performance. The user would be equally frustrated at home with an excellent connection. This is not the case with most airline passengers, who purchase inflight connectivity with an understanding of how the service should work.

This understanding can be defined. We make the following assertions in determining the performance model:

- The user is familiar with the application interface and how to interact with it.
- The user is generally aware of the time it takes to complete tasks within the application under typical conditions.

This knowledge forms the mental performance model the user brings to the inflight experience. The user likely is aware that there will be differences in application performance within the following use cases:

- Inflight Entertainment (IFE) (passenger-to-servers **in the aircraft**). The closest user analogy is at home on a wireline broadband connection.
- Inflight Communications (IFC) (passenger-to-servers **on the ground**). The closest user analogy is at wireless hotspots or in hotel rooms.

2.2. 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 grouped by two sets of user expectations for how the application interacts with the user.

- Request-reply applications: The application meets expected "responsiveness", e.g., the next screen starts to appear when expected.
- Continuous applications. The application delivers expected "flow", e.g., content continues to appear smoothly.

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

2.2.1. Request-reply Group

- 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]*
- Virtual Desk: Microsoft Office 365, Google G-Suite, corporate VPN access, remote SAP, etc. (aka virtual desk infrastructure (VDI))
- Gaming: multi-player real-time games like Fortnite, etc.

2.2.2. Continuous Group

- 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.

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.

2.3. Unsatisfactory Performance

This following are descriptions of some typical symptoms of poor performance that occur due to degraded communications performance. It is not intended to be an exhaustive list.

- Messaging: slow message sent confirmation
- Interactive: slow page load time
- VPN: dropped sessions, slow session recovery
- Virtual Desk: slow screen reaction time, mouse tracking, keyboard-screen reflection

- Gaming: slow player movement, high lag, other gamers 'beat you out', hard to stay alive
- File Transfer: transfer pauses, slow file load/transfer times, file corruption
- Streaming: video pixilation or loading interruptions, slow replay, slow restart from pause
- Social: pause in the scroll to more content, or video streaming issues like above
- Realtime: garbled voice, frozen video, pixilation, poor movement representation

2.4. Mapping Application Categories to Use Cases

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. The table below maps the applications to the use cases. It also serves as a guide for which applications may be enabled in the future as communications technology improves.

			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 – If use case has relevance.

2.5. 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.

2.5.1. Availability

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. This is the first metric that must be measured because if the service is unavailable, none of the other metrics matter or can even be measured. As addressed in **Section 1.1.5 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.

2.5.2. Response Time – Responsiveness

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.

2.5.2.1. Latency

Defined as round-trip time (RTT), 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 RTT 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. Each Get and reply is an RTT. A typical web page requires dozens of such RTT events.

2.5.2.2. Loss

Defined as a packet that is lost during an RTT event or TCP flow. 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. This is a second RTT event. 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,

2.5.2.3. Jitter

Jitter is the variation in one-way network transit time. Some applications such as multicast depend upon predictable transit time, which is separate from latency.

2.5.3. DNS Lookup

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.

2.5.4. Bandwidth

Bandwidth is defined as the effective number of bits per second (bps) or Bytes per second (Bps) successfully transferred from sender to receiver. Bandwidth has different benefits for single and multiple users.

If an application response requires a large amount of data, more bandwidth (i.e., connection capacity) will enable the content to move faster. However, as Figure 4 shows, some application categories require more time than others.

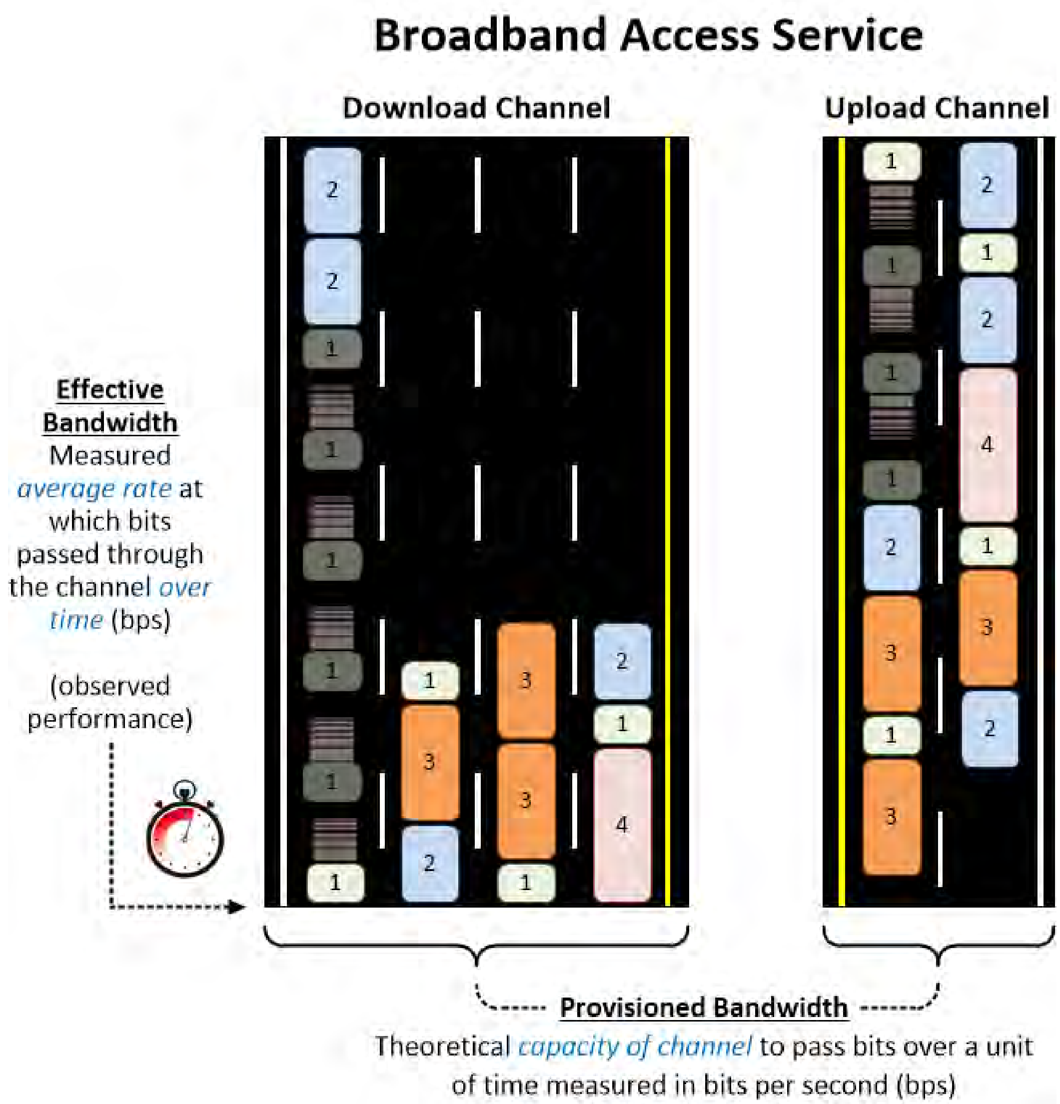

Figure 4. Provisioned versus Effective Bandwidth

Notes:

1. Numbered blocks represent different sized data payloads for different application types.
2. Provisioned bandwidth is typical of an asymmetric service where more capacity is allocated to data downloads.

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 the 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 any one 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 by the QoS Working Group.

2.5.4.1. Bandwidth Measurement

Provisioned bandwidth cannot be effectively measured from the seat. 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.

2.6. 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.

			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	Request-Reply	Messaging	Neutral	Neutral	Neutral	Low	Low
		Interactive	High	High	Neutral	High	Low
		VPN	High	High	Neutral	High	Low
		Virtual Desk	High	High	Neutral	Neutral	Low
		Gaming	High	High	Neutral	Neutral	High
	Continuous	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

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?

2.7. Evaluating Passenger Quality of Experience (QoE)

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

2.7.1. Testing the Metrics

True QoE measurement must be made via active tests. In order 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 to 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.

2.7.2. 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 communication 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).

2.7.3. 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 don’t interfere with the device use.

2.7.4. 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.

3. 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, 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 which 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-impacting 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 seat level as described in this document.

4. 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.
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.
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.** Personal 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.