SHARED COCKPIT

Quick Start Guide

Version 1.0 (2022-Dec-19)

Hot Start

CHALLENGER 650

SHARED COCKPIT

Feature Overview

Starting in version 1.7, the Hot Start Challenger 650 includes support for multi-crew operations with two users, fulfilling the roles of pilot and copilot. This feature works by running two copies of the CL650 on two separate computers, connecting them over the Internet, and synchronizing every aspect of appearance and state of the flight. To the users, this function is transparent - it simply appears that there is one aircraft, with two pilots inside, each pilot being able to operate all controls and functions of the airplane.

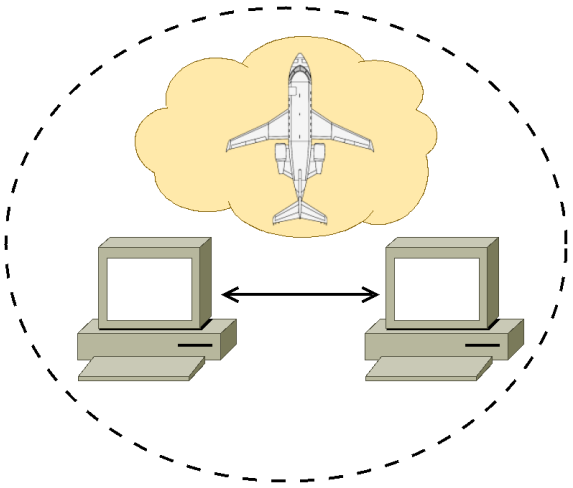

Prerequisites

To operate the CL650 shared cockpit feature, the following prerequisites must be met:

- Each pilot must have purchased their own copy of the CL650 and X-Plane. **Account sharing is strictly forbidden.**
- Match versions of the aircraft and operating system according to this table:

Pilot 1	Pilot 2
Windows 10/11	Windows 10/11
macOS or Linux	macOS or Linux

- Sufficient network bandwidth of at least 2 Mbit/s upload (on the host) and 2 Mbit/s download (on the guest). See below for details on which side is the host and which the guest.
- Mixing different X-Plane versions is discouraged, due to differences in scenery.

Principles of Operation

During operation, the computer of one user serves as the “host” for the shared cockpit session, while the other computer of the other user, known as a “guest”, connects to the host’s computer.

On the host’s side, the simulator runs essentially as normal. The host can select either Career or Non-persistent mode on the initial CL650 startup screen. From the point of view of the host, the shared cockpit feature operates as an option that sits passively in the background until a guest connects. The host can leave the shared cockpit feature enabled indefinitely, even when not expecting to fly with another person.

The guest needs to select shared cockpit operation from the initial aircraft startup screen. Upon connection, the guest's simulator will move to the exact location of the host, and their aircraft will mirror the exact state of the host's aircraft.

Technically, there are not actually two aircraft operating at the same time. Instead, the aircraft is entirely on the host's computer. On the guest's computer, there is only a minimal set of modules active, primarily the ones for the flight displays and sound generation. The simulator on the guest also forwards any user inputs (flight control movements, keyboard bindings, etc.) to the host, where they are processed, and the results are then sent back to the guest for display. As such, there's almost no possibility of desynchronization between the two sides, since the guest doesn't actually have a full aircraft operating. All the systems, physics, flight management, and avionics are only present on the host, with the guest being able to also interact with those systems remotely.

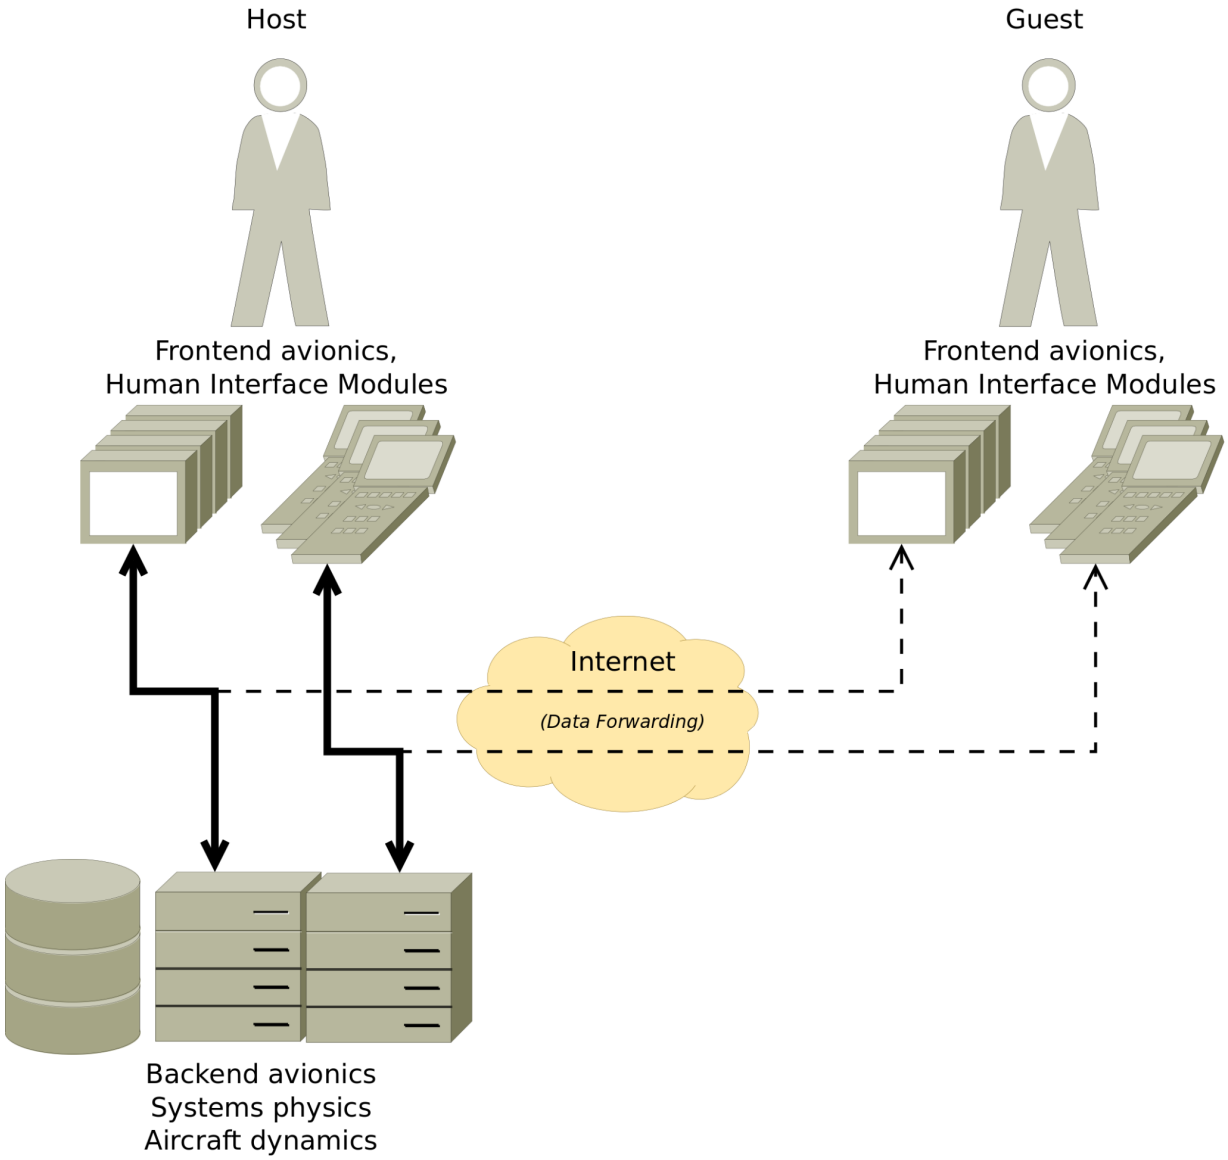

The practical implications of this technical arrangement are as follows:

- The guest's simulator environment is entirely irrelevant to the physical state of the aircraft. The guest is encouraged to configure their weather identically to the host (using either real weather, manually set weather, or a 3rd party weather-injecting plugin), but this is only to facilitate a similar visual presentation to the guest.

- Environmental factors, such as the location of visible moisture, precipitation and cloud buildups may be difficult for the guest to judge accurately, since X-Plane weather is non-deterministic, meaning even if both users are using synchronized “real” weather, the exact weather depiction may differ between them.
- The navigational database state only matters on the host. The guest’s navigational database isn’t used for any functions of the aircraft avionics and systems simulation.
- Both users should be using the same scenery (whether customized or stock). Otherwise, visual differences in elements such as hangar building placement, taxiway layout or airport markings can lead to confusion between the pilots.
- While the shared cockpit feature attempts to account for terrain height differences by utilizing a height blending model, this can lead to seemingly unnatural flight behavior of the aircraft to the guest. This blending starts at 500 ft AGL and gradually as the aircraft approaches the ground on the host, the guest will start to blend true elevation with height above local terrain. If there are significant terrain elevation differences between host and guest, the aircraft can appear to be suddenly rising or sinking, to make sure it makes ground contact at the same point in time for both host and guest.

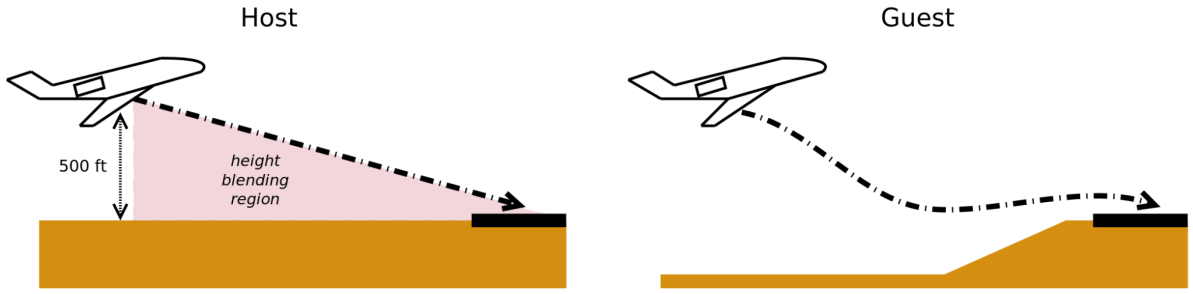

- While the aircraft can be hand flown by either host or guest, network transmission time delay can affect the feel and responsiveness to control inputs for the guest. This should be taken into account by the guest user, to make sure they don’t overcontrol or enter into a pilot-induced oscillation due to excessive and rapid corrections. See the section Operational Considerations for more information.

Host Network Configuration

Since the host computer will often be sitting behind a firewalling router, in order to be able to receive the incoming guest connection, some configuration may be required. This typically takes the form of some kind of network port forwarding. Using this method effectively “punches a hole” in the firewall, but only for the purposes of forwarding the required connection from the guest to the host, rather than fully exposing the host to any arbitrary outside network traffic.

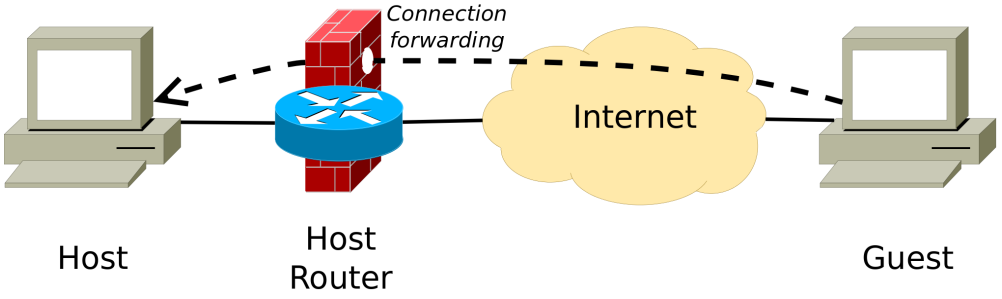

There are a three ways of achieving this kind of network forwarding:

Automatic Configuration Using UPnP

If the host’s router has the UPnP protocol enabled (or the user enables it in the router’s configuration), then the port forward is set up automatically whenever the CL650 is loaded in the simulator and shared cockpit hosting mode has been enabled in the aircraft’s User Settings. This option requires minimal user involvement and is the most convenient. If UPnP auto-configuration succeeds, you will see the following window appear when you load the CL650 with shared cockpit hosting enabled:

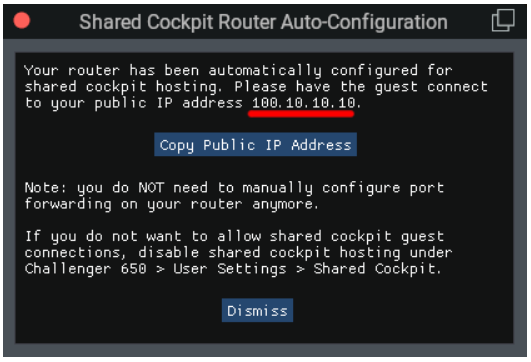

This confirms that the aircraft has been able to communicate with the network router and set up the port forward from the router’s public IP address. The highlighted address listed in this window is the “Host Address” that the guest needs to enter during connection. Clicking on “Copy Public IP Address” copies the listed IP address into the clipboard, so you can more easily forward the address to your shared cockpit partner.
This window will reappear whenever your public IP address changes, to remind you that you have shared cockpit hosting enabled.

Manual Configuration on the Network Router

If network auto-configuration doesn’t work, an alternative is to attempt manual configuration of the port forward. For that you will need administrative access to your network router. Unfortunately,

every manufacturer styles their router configuration interface differently, so we cannot provide an exact step-by-step guide.

Before we begin, we need to determine the host computer's internal IP address.

- Windows 10/11:
    a. Right-click on the network icon in the status bar and select "Open Network & Internet settings".
    b. Under "Network Status" and "Ethernet" (or "Wi-Fi"), select "Properties".
    c. Scroll to the "Properties" section and copy the IPv4 address listed as shown on the right.

- macOS:
    a. Click the Apple logo on the menu bar and select "System Preferences".
    b. Select "Network".
    c. Select your main network connection from the side menu and copy the IP address as shown on the right.

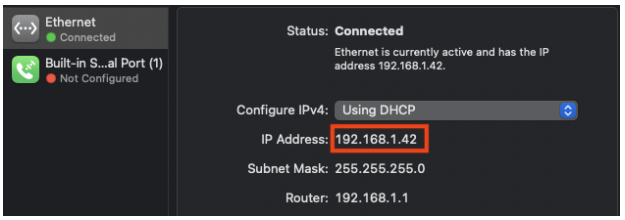

- Linux:
    a. Open a terminal window and enter the command "ifconfig".
    b. Copy the IP address for your main network interface from the field labeled "inet" as shown on the right. Please keep in mind that multiple interfaces may be shown. The Ethernet interface will typically be named "en<*something*>" and have an IP address in the form of 10.x.x.x or 192.168.x.x.

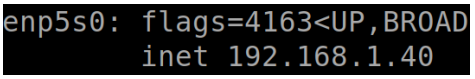

With the computer's internal IP address in hand, we can configure the port forward. Verify on your router that your external (or "WAN") IP address is indeed public. This will typically be visible in a section labeled "Status" and the external interface is usually named "WAN". Verify that the WAN IP address **DOESN'T** match any one of the following formats (the "x" in the listings below indicate that any number may be present in that part of the address):

- 10.x.x.x
- 172.16.x.x - 172.32.x.x
- 192.168.x.x

These network addresses aren't public and cannot be used by your shared cockpit partner to initiate a connection. If your WAN IP address matches one of these, you will need to fall back to the last method, listed in the section Virtual Private Network Using Hamachi.

If your WAN IP address **doesn't** match any of the above formats, it is public and should be usable for manual port forwarding. As noted before, the exact configuration menus will differ based on router manufacturer, so we cannot provide an exact step-by-step guide. However, the general gist is as follows:

- Navigate to your router's port forwarding configuration. This can be named "Virtual Servers", "NAT forwarding", "Port forwarding" or some other combination of these terms.
- Create a new port forwarding rule by clicking a button named "Add" or "+".

- Enter a name for the rule. This is simply something descriptive for you.
- Under “External port” or “WAN port” enter 8205.
- Under “Internal IP” or “LAN IP” enter the IP address of your computer that we determined at the beginning of this section (e.g. “192.168.0.50”).
- Under “Internal port” or “LAN port” enter 8205.
- Under “Protocol” select ALL. If the only selectable options are TCP and UDP, create two port forwarding rules, one for “TCP” and one for “UDP”. Both rules are required for operation.
- Click Save.

See below for an example network configuration interface on a TP-Link router:

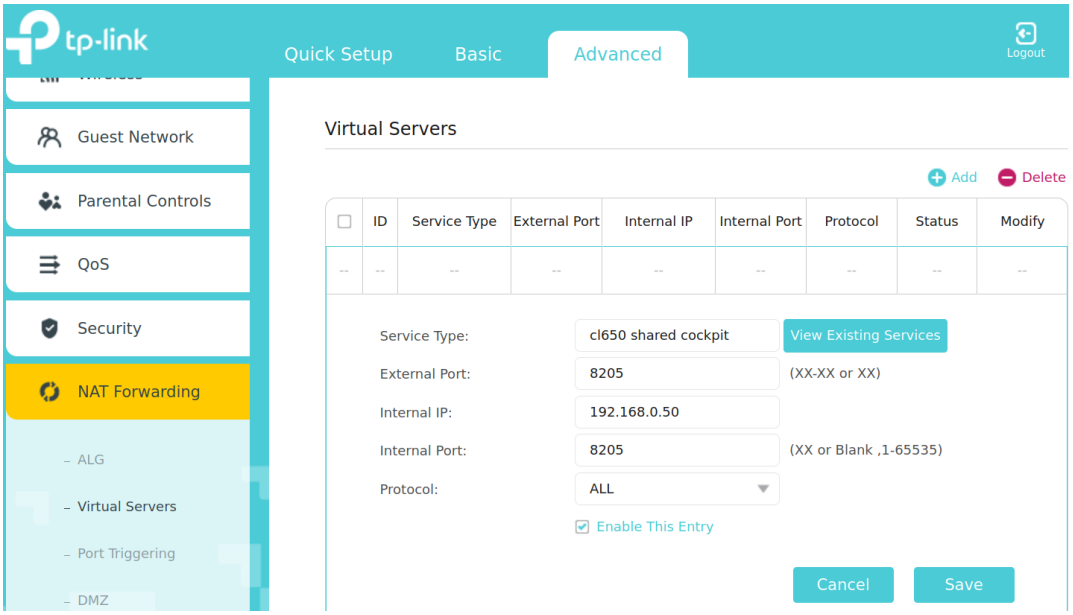

When saved, you will need to send your public WAN IP address to your shared cockpit partner, as that is the “Host Address” that they will need to enter during connection.

Virtual Private Network Using Hamachi

Hamachi by LogMeIn is a VPN network auto-configuration tool that can be used on Windows 10/11, macOS and Linux to create a private network between the host and guest. Please consult the Hamachi user guide on its usage.

Once connected via Hamachi, the guest will enter the host’s Hamachi private IP address into the Host Address field during connection.

Host Simulator Configuration

Once the network is configured, the host can enable shared cockpit operation. To do so, simply load into the simulator with the CL650 and select either Career or Non-persistent mode. Once loaded, from the simulator menu select Challenger 650 > User Settings… In the User Settings window, navigate to the “Shared Cockpit” tab:

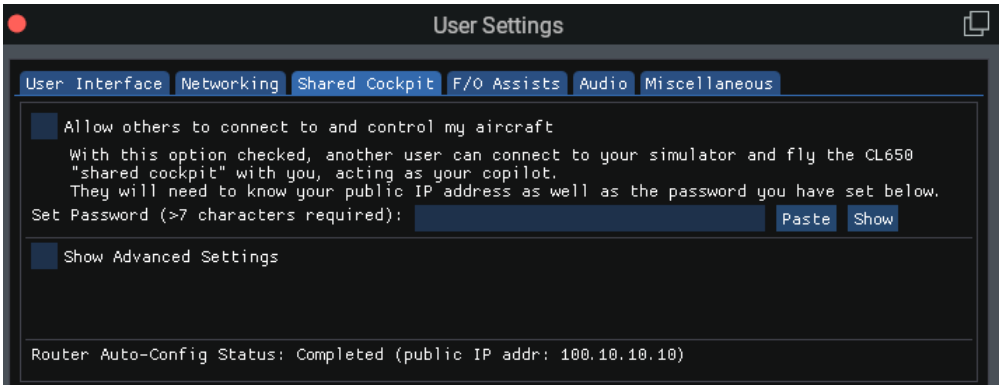

Before enabling shared cockpit hosting for the first time, you will need to set an access password and communicate this password to your shared cockpit partner. This serves as an extra security mechanism, to prevent unauthorized shared cockpit access. Once set, check the checkbox labeled “Allow others to connect to and control my aircraft”. Before commencing shared cockpit hosting, a soft reload of the aircraft is required. You can accomplish this by clicking the “Reload Aircraft” button which appears after checking the above checkbox. This may take a few seconds and the simulator may appear to be “hung” during that time, please be patient.

With these steps completed, the host is now fully configured and ready for shared cockpit operation. When you start X-Plane the next time, no further configuration actions are required. The entire configuration is memorized and the guest can connect right away.

Guest Connection Initiation

The guest can connect by selecting “Network Guest Mode” on the CL650 startup screen. This will show the following screen:

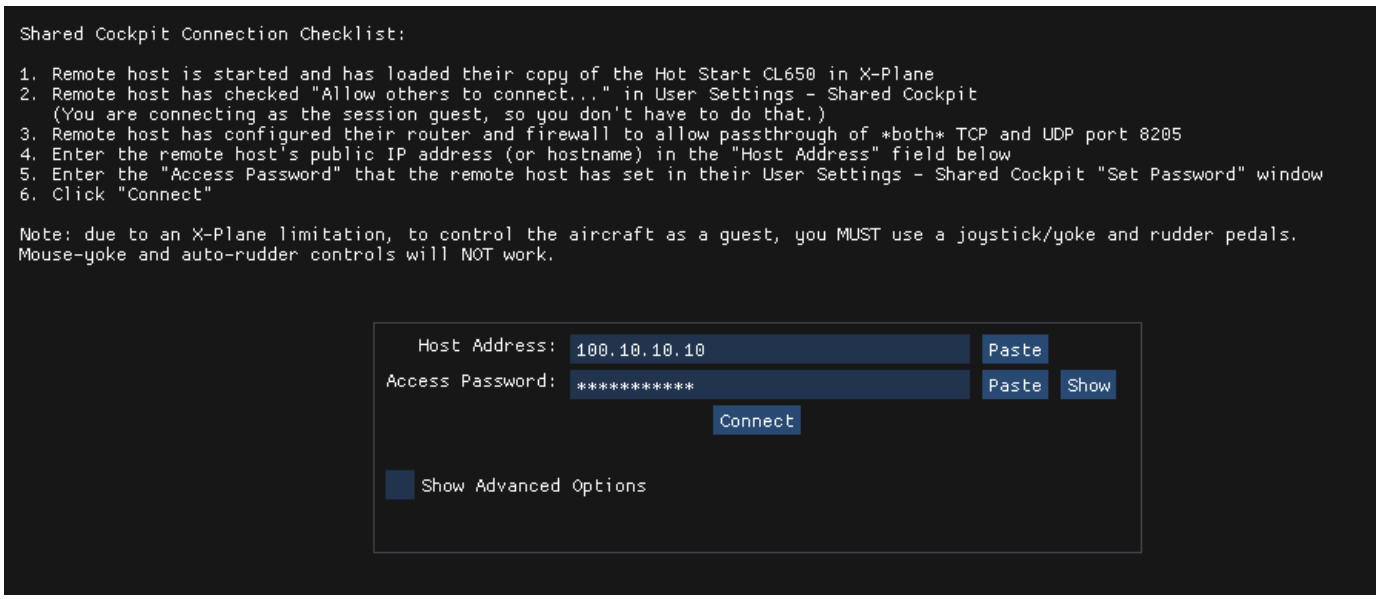

Prior to attempting connection, verify the actions on the Shared Cockpit Connection Checklist have been accomplished.

Enter the “Host Address” and “Access Password” as communicated by the host, then click “Connect”. During connection setup, the host will be prompted to confirm the incoming connection request as shown below.

- To allow the connection request, select “>> Allow <<”. This permits connections from the guest’s IP address for the remainder of the simulator session and won’t prompt you again.
- To reject the connection request, select “Reject”. This blocks reconnection for a few seconds and immediately displays the following message to the guest:
  “Cannot connect to <address>: host rejected our connection request (wait at least 5 seconds before retrying)”
- If the incoming connection attempts are malicious and/or abusive, select “Reject & Block”. Any further connection attempts from this IP address will be silently rejected for the remainder of the simulator session. If you selected this option by mistake, you will need to restart the simulator to allow the guest to connect.

Connection Troubleshooting

If a connection failure occurs, you will see one of these messages:

- *“Cannot connect to <address>: Connection refused”*
  This typically happens due to X-Plane not being started on the host machine, or a firewall rule exists that immediately rejects all incoming connection attempts.
- *“Cannot connect to <address>: Timed out”*
  This occurs when there is a mistake in the Host Address, the host’s port forwarding configuration doesn’t work, or the host is running a restrictive firewall that ignores the incoming connection attempt.
- *“Cannot connect to <address>: host rejected our connection request”*
  This indicates that the host user has explicitly rejected the guest’s connection attempt.
- *“Cannot send message to <address>: <reason>”*
  This is a special corner case that can happen if the guest’s computer doesn’t allow outbound traffic. Verify that the guest computer firewall isn’t set up too strictly, or disable it.

Pilot Roles

After connection, the guest automatically assumes the role of the right-seater (copilot). This isn’t fixed, however. If a seat swap is required, it can be performed while the aircraft is either stationary on the ground, or in flight, when at a height of greater than 1,300 ft AGL. To request to swap seats either:

- Select “Challenger 650” > “Shared Cockpit” > “Request to Swap Seats” from the simulator’s main menu; or,
- Click on the seat cushion of the other pilot.

When a seat swap is requested, the other pilot must agree to it by selecting “Accept” from the window shown on the right.

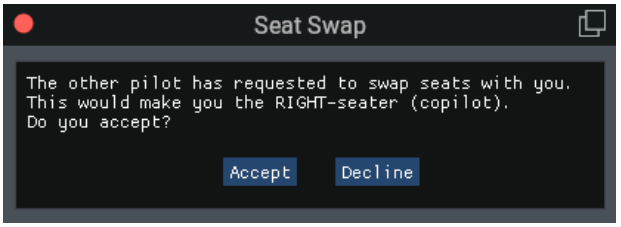

The currently active role is indicated using a status bar at the top of the screen. It shows “L” on the screen of the currently active left-seater and “R” on the screen of the right-seater:

Controls

Either pilot can make control inputs at any time. There is no explicit handover of controls in the software. Therefore, correct cockpit CRM should be observed, handing over control by standard verbal callouts such as “I have control” and “you have control.” Also, in shared cockpit mode, control inputs no longer follow the view position of the user. They are instead “pinned” to the respective control side, depending on the assumed role - left-seater always to the left controls, right-seater always to the right controls.

Control inputs on the roll & pitch axis follow these rules:

- While the respective axis disconnect handle is stowed, inputs from the two pilots are summed up. You can think of this as the two pilots applying force to the control, and the resulting control position being the sum of the mechanical forces. For example, if both pilots are giving a 50% right roll input, the result is a 100% right roll command. Conversely, if one pilot is giving a 50% left roll input, and the other pilot a 50% right roll input, the result is a zero roll command (control wheel neutral).
- When the control axis disconnect handle is pulled, inputs from each of the pilots are simply applied directly to their respective control wheel and/or column.

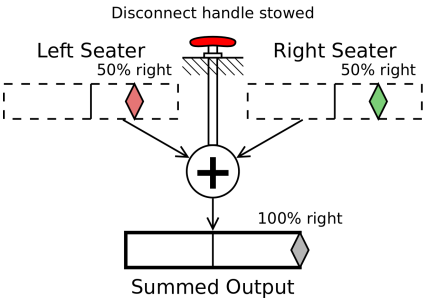

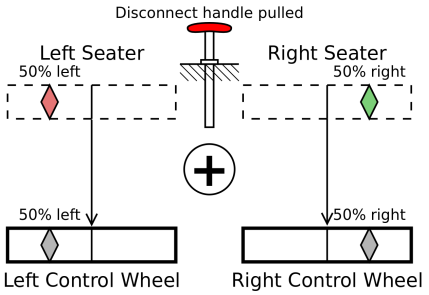

Control inputs on the yaw axis are always summed up between the two pilots, since the yaw axis inputs are independent and only synchronized through springs that can be overcome with force.

Throttle input is somewhat special, since there's only one set of throttles in the aircraft. The last crew member who made a significant throttle input (>5% axis movement) will have exclusive control of the throttle. If one of the pilots has a very noisy axis input which causes frequent throttle control takeover, they should recalibrate or disconnect their respective throttle controls, to avoid interfering with the pilot flying.

Please also note that the right-seater does **not** have control of the steering tiller, so taxiing of the aircraft can only be done by the left-seater.

Interactions with Ground Handling

In Career mode, either crew member can enter and interact with the FBO. Typically it'll be the responsibility of the fight's commander to submit a handling request, but the copilot can do so as well. When calling to inform the crew of the passengers having arrived at the FBO, Jenny will only call the flight's commander (left-seater). However, outgoing calls to the FBO and the passengers can be made by either crew member, using either the cell phone, or satellite phone.

When the fueler arrives and enters the aircraft, he will first look for the commander (left-seater). If that person is not present inside the aircraft at the time that the fueler attempts to enter, he will try to speak to the copilot (right-seater), if present. If neither crew member is currently inside the aircraft, the fueler will wait for one of them to enter the aircraft.

Either crew member can interact with the de-ice truck. To do so, the crew member wanting to place a call to the de-ice operator should tune to 136.925 MHz on their respective radio set and properly set up their audio control panel, to be able to speak and receive transmissions.

Flying with Online ATC

Online flying is fully supported in shared cockpit mode on both Vatsim and PilotEdge. Each pilot's radio reception and transmission is configured according to their own respective audio control panel settings. This will follow the left/right seat roles, as currently assigned.

Shared cockpit online flying on IVAO or POSCON has not been tested.

PilotEdge

One of the crew members, typically the host, will connect to PilotEdge as normal. The second crew member should then connect using the same callsign with the '@' character appended to the end of the callsign. No further special configuration is required. The PilotEdge server system prevents loopback of crew audio transmissions back to the other crew member, so while other users on the network can hear each crew member speaking, the crew members will not be hearing each other through their radios.

Vatsim (xPilot client)

One of the crew members, typically the host, will connect to Vatsim as normal. The second crew member should then connect using the same callsign, but add a single letter (A–Z) at the end of the callsign to differentiate the connection, and also check "Connect in Shared Cockpit/Observer Mode". For example, if pilot 1 connected as "N123AB", then pilot 2 should connect as "N123ABA" and check the box for shared cockpit.

Unfortunately Vatsim doesn't properly implement audio loopback prevention. If pilot 1 transmits while pilot 2 is monitoring the same frequency (either on the same radio, or on another radio), Vatsim doesn't properly inhibit the audio being retransmitted to pilot 2. If the pilots are also using some kind of private voice chat system such as Discord, this Vatsim audio loopback results in doubled audio being heard by pilot 2 (once through Discord and once through Vatsim). The suggested workaround is to enable push-to-mute on Discord for each pilot. That way, whenever they transmit on the radios, they will be muted on Discord.

Operational Considerations

This section lists various operational considerations and aspects of the shared cockpit function, in no particular order:

- The performance impact of shared cockpit operation is negligible, although additional network traffic will occur. The typical data streaming load is as follows:

Side	Upload	Download
Host	800 kbit/s	100 kbit/s
Guest	100 kbit/s	800 kbit/s

- Since user inputs are first sent from the guest to the host, where they are implemented, after which the resulting effects are sent to the guest, network latency can affect the control feel for the guest while hand-flying. The guest continuously monitors network latency and displays it in the status window at the top of the screen, as depicted on the right.

The following table lists approximate control feel for the guest at various latency values:

Latency	Rating	Perceived control response
< 50ms	Excellent	Essentially imperceptible
50ms - 125ms	Good	Slightly perceptible, normal flying technique unaffected
125ms - 200ms	Tolerable	Perceptible, requires gentler inputs to avoid PIO in landing flare
> 200ms	Poor	Very perceptible, difficulty with control in turbulence or in the landing flare

Hand-flying shouldn't be attempted by the guest when network latency exceeds 200ms.

- When the host pauses (explicitly or by opening an X-Plane menu), or enters replay mode, the guest's simulator pauses and waits for the host to resume normal flight.
- When the connection to the host is lost, the guest pauses and waits for the host connection to become available. In case the host simulator crashes, the host can restart the simulator and resume the last state save. The guest simulator will automatically reconnect and resume flight from the last saved state.
- The Failure Manager is available for either pilot and is fully synchronized. This facilitates training, where one of the pilots can serve the role of training instructor.

Document Revision History

Version	Date	Changes
1.0	2022-Dec-19	Initial release of document