Pathfinder Intel Recon

recon.pathfinderil2.com

The objective of Pathfinder Intel Recon is to produce, as automatically as possible, annotated reconnaissance photographs of targets for IL-2 Great Battles.

Pathfinder has the ability to see virtually everything contained within a mission file, but it is deliberately designed not to show everything. Instead, it attempts to reproduce how a wartime Photo Interpreter (PI) might have assessed a reconnaissance photograph, retaining uncertainty, interpretation and an appropriate degree of fog of war.

Step 1 — Prepare the Mission

The starting point is the IL-2 `.MISSION` file.

The mission is first passed through a preparation stage where:

- Weather conditions are normalised to clear skies.
- Time of day is moved to early afternoon to provide consistent lighting.
- Mission Editor markers, trigger points and other objects visible in the IL-2 Mission Viewer but irrelevant to reconnaissance photography are removed.

This produces a clean, recon-ready version of the mission.

Step 2 — Pathfinder Intel Logic Engine

The prepared mission is then passed through the Pathfinder Intel logic engine.

At its core is a workbook that maps the large number of IL-2 object types into approximately 15 broader Pathfinder classifications.

For example:

- A Bf110 may be classified simply as a Twin-Engine Aircraft.
- Different models of tanks may all be classified as Armour/Tanks.
- Different hangar models can be grouped by size and operational importance.

For each Pathfinder classification, the workbook contains information such as:

- Physical size.
- Vulnerability.
- Likelihood of being hit by different types of ordnance.
- Typical attack priority.
- Whether the object normally appears individually, in isolated clusters or in tight groups.
- Relative operational or supply importance.

The workbook also contains a lexicon of alternative PI terminology so that annotations and intelligence reports do not repeatedly use identical wording.

For example, armour might variously be described as:

- Tanks
- Heavy tanks
- Tracked vehicles
- Armour

Damage descriptions can similarly vary:

- Burning hulks
- Turrets destroyed
- Tracks shed
- Vehicles knocked out

Terminology can also change according to context. For example, a fuel storage object at an armour encampment might simply be described as

fuel stores, while the same object at an airfield might be described as AVGAS.

Further hard-coded rules sit alongside the workbook to refine the interpretation. For example, a bridge may be disregarded as an objective unless other significant objects — such as bunkers, AA positions or military installations — are located nearby.

The `.MISSION` file provides the exact spatial relationships between objects. Pathfinder then uses the classification data to determine whether those objects represent:

- Individual targets.
- Tight target clusters.
- Long repeating or linear targets.
- Dispersed high-value targets.
- Background or irrelevant scenery.

Once the target structure has been mapped, Pathfinder determines which objectives should actually be photographed, with priority influenced by their operational and supply importance.

Where possible, Pathfinder attempts to capture an objective in a single reconnaissance photograph. If the target is too large or dispersed, it calculates the number of photographs required.

Step 3 — Automated Recon Photography

Pathfinder then takes control of the IL-2 Mission Editor camera.

The camera is normally moved into an orthographic overhead view, although experimental angled photography is also working well for some compact targets.

For each view, Pathfinder captures four images, which are tiled together to produce a single high-resolution reconnaissance photograph.

The system normally prefers a north-up orientation, but it can rotate the photograph where this produces a better presentation — particularly for predominantly linear objectives such as rail yards, roads, docks or elongated industrial sites.

It can also disregard small outlying groups when including them would force the main target to be photographed at an unnecessarily low scale.

Further work is still required around the automatic padding applied to targets, as there are currently some inconsistencies in how the optimum framing is calculated.

For very large or widely dispersed objectives, Pathfinder takes as many photographs as necessary and records the spatial relationship between them in the associated JSON data.

A target currently takes approximately one to two minutes to acquire. Much of this time is spent finding the correct Mission Editor zoom level by emulating mouse-wheel movements. The IL-2 zoom behaviour is non-linear, and so far I have not found a more reliable method of controlling it.

Additional render time is also allowed to ensure that all mission objects have appeared correctly before the images are captured.

On the current system — RTX 5090, 32 GB RAM and AMD X3D CPU — a typical reconnaissance collection run takes approximately 30–60 minutes, depending on the mission.

Individual objectives can also be acquired manually if required.

Step 4 — Target JSON

Using the intelligence engine, mission data and IL-2 map information, Pathfinder generates a JSON file for each photographed objective.

Typical information includes:

- Objective name.
- Location.
- Coalition.
- Target objects and their locations.
- Target numbers and relative importance.
- Nearest alternative objectives within approximately 50 km.
- Nearest airfields.
- Rotation from north.
- Elevation at the photographed objective.
- Size of the photographed area.
- Nearby towns and landmarks.
- Numerical terrain information describing hills, valleys and surrounding ground.

Coalition information stored in the static recon JSON is essentially the underlying identification. Where a live campaign is running, current ownership can subsequently be applied as a separate campaign-state layer.

Step 5 — Image Processing and PI Annotation

The captured PNG and its target JSON are then passed through the Pathfinder image-processing and annotation system.

A number of processing options are available, including:

- Colour or black-and-white output.
- Film grain and photographic ageing.
- Negative scratching.

- Bloom.
- Light or heavy photographic marking.
- Period stamps and administrative marks.
- Different output types for different uses.

Standard information can then be added to the image.

Top left

- Objective name.
- Target type.
- Grid reference.
- Elevation.

Bottom left

- Photograph scale.

Bottom right

- Reconnaissance information strip.
- Capture altitude and related photographic information.
- North arrow.

Where an objective requires more than one photograph, the images can also show their spatial relationship to one another.

PI annotations

Annotation density can be selected from none through to full annotation.

A heavily annotated photograph may contain up to approximately six marked areas.

Pathfinder knows the relative priority of the identified targets and attempts to:

- Avoid overlapping annotations.
- Avoid obscuring important target detail.
- Prefer period-style ovals around target groups.
- Use arrows for discrete high-value targets.
- Use close-fitting or "hugging" lines around some larger groups.

- Route annotation lines around other labels where possible.

The terminology selected for a target remains consistent within an individual photograph but may change between photographs, reproducing some of the natural variation expected from different PI reports.

Several handwriting styles are available, with one style normally remaining consistent across an individual image.

AAA / Flak

AAA is divided into three broad classifications:

- Light Flak.
- Heavy Flak.
- Puffy/airborne Flak.

The level of AAA annotation can be selected independently:

Off / 25% / 75% / Full

The default is Off.

For dynamic campaigns such as Apollo, showing every Flak position on the underlying reconnaissance photograph would be misleading because those positions can change from day to day. Static photographs therefore normally omit them, although the option remains useful for conventional or curated missions.

Step 6 — Cloudflare Upload

The finished PNG files and their associated JSON files are currently manually uploaded to Cloudflare.

This is one of the remaining stages that could potentially be automated.

Step 7 — Pathfinder Recon Website

Cloudflare presents the reconnaissance material through the Pathfinder Recon website in two main formats:

- Flat-screen viewing.
- VR-optimised viewing through OpenKneeboard Web Capture.

For servers such as Combat Box that expose suitable JSON or API data, Pathfinder can create Smart Links.

A Smart Link is effectively a permanent URL that automatically resolves to the reconnaissance material for the mission, theatre or campaign currently in play.

This means a pilot does not need a different kneeboard URL every time the server rotates to a new mission.

Step 8 — Live Intelligence Layer

Where live server data is available — currently including Combat Box — Pathfinder can combine the static reconnaissance photograph with information describing the mission as it exists now.

Weather was originally presented using a period Manila-style card over the photograph. Testing, particularly in VR, showed that this consumed too much image space and reduced readability. The current design therefore places supporting intelligence above and below the main reconnaissance photograph, leaving the photograph itself as unobstructed as possible.

Because each reconnaissance JSON also records nearby objectives and airfields, Pathfinder can provide additional contextual information such as:

- Current weather.
- Nearby objectives.
- Alternative targets.
- Friendly airfields.
- RTB headings.

- Local terrain considerations.
- Current coalition ownership, where live campaign data is available.

Each objective also has information describing its target types, locations, groupings, preferred attack direction, photographic orientation, elevation and surrounding terrain.

This can then be combined with current supply/damage status and Pathfinder Intel rules to generate a dynamic, period-style intelligence assessment for that particular objective on that particular campaign day.

The intention is not simply to display numbers such as "Target 63% damaged". Pathfinder instead turns those numbers into an immersive PI-style assessment — for example describing damaged hangars, burning aircraft, disrupted fuel storage or surviving concentrations of armour according to what the underlying campaign data implies.

Two versions of the assessment can be produced:

Long Recon Analysis — automatically displayed initially for approximately 45 seconds, giving the pilot time to read a fuller intelligence briefing.

It then fades away to reveal the reconnaissance photograph.

Short Recon Briefing — remains visible alongside the photograph, containing the most important current intelligence together with weather, nearby objectives, RTB information and other operationally useful details.

The result is therefore a combination of a static historical-style reconnaissance photograph and a dynamic intelligence layer representing the current mission or campaign state — while deliberately retaining enough uncertainty that Pathfinder provides intelligence rather than perfect knowledge.

A learning loop is built in whereby recon analysis is reviewed and improved constantly