

# HAPS Alliance

HIGH ALTITUDE PLATFORM STATION

## Acceptable Levels of Risk for HAPS

# Executive Summary

This paper discusses setting acceptable levels of risk for High Altitude Platform Systems (HAPS).

We first discuss how the safety metrics traditionally used in aviation (e.g., the probability of catastrophic accident per flight hour, or per mission) are not adequate to establish target levels of safety for HAPS. Instead, we propose to use “third-party-centric” metrics that measure the risk from the perspective of the exposed third parties (risk per human hour/year). We suggest a 2x2 matrix for acceptable risk levels that considers individual and collective risk limits for ground and air risk.

We then propose to set acceptable levels of risk to be consistent with the risks already accepted by the exposed parties and to look at examples from aviation and outside aviation, such as infrastructure risk to the uninvolved public. In particular, we draw inspiration from the UK HSE’s ALARP framework.

Finally, we propose a framework by which an operator self-manages the collective risk it generates through active control of population/aircraft density overflow, time spent in an area, the number of platforms used and the mix of platforms used. This approach would enable operators/manufacturers to safely operate sooner -- starting at low operational volume in low-risk regions, learning from flight experience and improving their platforms while generating the revenue necessary to sustain R&D, before expanding to higher-risk regions. This approach could also change the meaning of “certification,” which would no longer be a binary approval to operate commercially, but rather a certification that the “risk-rate” (value on a continuous scale), used in the accumulation of risk is accurate.

## Table of Contents

Executive Summary .....	2
Aviation Risk Metrics Do Not Work for HAPS .....	3
Third-Party-Centric Individual Risk.....	4
Societal (Collective) Risk for HAPS.....	7
Summary of Proposed Acceptable Levels of Risk for HAPS .....	9
Operator-Managed Societal Risk .....	9
Annex – Examples of Why Aviation Metrics Do Not Work .....	12
Example #1 – Per-flight-hour risk metric is inadequate .....	12
Example #2 - Per-mission risk metric is inadequate.....	13
Join Us in Our Work.....	14
About the HAPS Alliance .....	14



## Aviation Risk Metrics Do Not Work for HAPS


Traditionally, aviation has used safety metrics that measure the risk (typically a probability of catastrophic accident) on a per-flight-hour, or per-mission (per-flight) basis. Lin et al. 2009<sup>1</sup> give an excellent summary of historic aviation target levels of safety and societal expectations of risk.

Per-flight-hour or per-mission safety metrics work well for passenger transport because they measure the risk with a unit of time that relates directly to the exposed individuals (the people on board).

These metrics are, however, inadequate to quantify the risk for High Altitude Platform Systems (HAPS) which do NOT carry people on board. Per-flight-hour or per-mission metrics are “**platform-centric**” and can promote system designs that are misaligned with true safety goals. In particular, platform-centric metrics will disadvantage platforms with longer mission duration and will disadvantage larger platforms even when those create safer overall systems (see examples for details).

Because HAPS do not carry people, they create a risk to third parties exclusively. As they operate above commercial traffic, the main source of risk comes from the possibility of an “unplanned descent,” which can create the following two types of risk:

- Risk to populations on the ground, general public (third parties)
- Risk to manned air traffic operating below HAPS: mid-air collision during unplanned descent.(third parties)



**Figure 1** – HAPS operate above aircraft traffic. Unplanned descents are the main source of risk – risk to ground populations and risk of a mid-air collision with manned aircraft traffic operating below.

<sup>1</sup> Lin X, Fulton NL, *Target level of safety measures in air transportation – review, validation and recommendations*, Proceedings of the IASTED International Conference, Modelling, Simulation and Identification, October 12-14, Beijing, China, 2009

From a safety point of view, HAPS can be viewed as **airborne infrastructure**. They are networks of “flying cell phone towers,” earth-monitoring devices, or other kinds of semi-permanent infrastructure. In many cases, continuous coverage is maintained by cycling platforms during maintenance such that there is always a constant number of platforms in the sky (e.g., to ensure constant connectivity service). Even though the actual platforms and missions may be cycled, from the perspective of exposed populations, there is always the same number of platform(s) in the sky.

Therefore, to evaluate HAPS safety, we must use metrics and acceptable levels of risk that consider the system as a whole, and account for platform density. We cannot use platform-centric metrics (per-mission or per-flight-hour), because the number of missions (maintenance cycles) or number of flight hours (platform density) can vary greatly with system designs (see examples in the Annex).

HAPS risk is then per **real hour of the exposed party, or more practically per year, and it needs to encompass all HAPS in a region**, not just those from one operator.

This paper proposes safety metrics and acceptable levels of risk for HAPS that are **third-party-centric** (i.e., where the risk is measured from the perspective of the ground-individual-exposed aircraft operating below the HAPS fleet) rather than the traditional aviation metrics that would measure the risk from the platform's perspective

In proposing acceptable levels of risks, we consider both the **individual risk** (the risk to each individual exposed) as well as a **collective risk or societal risk** (the risk to a group of people).

We use **comparable risks** to propose acceptable levels of risk that are in line with risk levels currently accepted by the exposed parties:

- Aviation risk standards to establish an acceptable level of risk to exposed aircraft
- Infrastructure standards (such as the UK HSE ALARP Framework<sup>2</sup> for pipelines, power plants, dams, industrial plants, etc.) to establish the risk for ground populations

Finally, we propose a **dynamic collective risk management** in which each HAPS operator computes the collective risk integral (using actual trajectories, population density data and aircraft density data) and is responsible for maintaining the collective risk under the defined acceptable level of risk.

The approach proposed in this paper is consistent with the work performed by the ICAO Separation Airspace Safety Panel (SASP) for HAPS<sup>3</sup>.

## Third-Party-Centric Individual Risk

To appropriately evaluate the risk created by HAPS, and incentivize appropriate safety decisions, we must consider a system as a whole, rather than each platform (or aircraft) individually.

---

<sup>2</sup> [https://www.hse.gov.uk/foi/internalops/hid\\_circs/permissioning/spc\\_perm\\_37/#Tools-for-ALARP](https://www.hse.gov.uk/foi/internalops/hid_circs/permissioning/spc_perm_37/#Tools-for-ALARP)

<sup>3</sup> SI Barry, High Altitude Platform Systems: guidance material: Attachment on mathematical Modelling, ICAO Separation Airspace Safety Panel, Working Paper 4, Montreal, 19 May, 2023

An adequate risk metric is one that measures the risk per unit of time of an exposed individual or exposed aircraft (third-party-centric), rather than a unit of time associated with the uncrewed platform (platform-centric). We propose the following:

- The risk to manned aircraft flying below HAPS should be measured as the probability of mid-air collision per exposed manned aircraft flying hour.
- The risk to ground populations living in the HAPS service area should be measured as the probability of being fatally impacted per exposed person per year.

Note that the above metrics intrinsically account for the density of HAPS.

Because the risk is measured in the exposed party's frame of reference, we can set the acceptable level of risk to match other risks already accepted by the exposed party:

- **For mid-air collision** -- ICAO uses a target level of safety of  $1.5 \times 10^{-8}$  per aircraft flight hour for en-route separation. Because the HAPS risk is additional, the ICAO SASP has used  $5 \times 10^{-9}$  per exposed (manned) aircraft flight hour in its work on unmanned free balloons, which is currently being generalized to HAPS<sup>3</sup>.

We propose to keep this acceptable level of risk of  $5 \times 10^{-9}$  mid-air collision per exposed manned aircraft flight hour. (Note: we may want to consider setting a separate risk for general aviation, which typically accepts a higher risk.)

- **For ground populations** -- Because HAPS are adding to the risk experienced by ground populations, it is relevant to set the acceptable level of risk from HAPS in line with the risks that other infrastructures create on ground populations.
  - The UK HSE ALARP<sup>4</sup> Framework sets the acceptable level of risk to the general public for pipelines, power plants, etc. between  $10^{-4}$  (limit of tolerable risk) and  $10^{-6}$  (limit of broadly acceptable risk) per person per year. For example, the likelihood for a pedestrian to be fatally struck by a car in the United States is  $2.2 \times 10^{-5}$  each year<sup>5</sup>.

Consistent with the UK ALARP, we propose to set a tolerable range between  $10^{-4}$  and  $10^{-6}$  probability of fatality per person per year.

---

<sup>4</sup> [https://www.hse.gov.uk/foi/internalops/hid\\_circs/permissioning/spc\\_perm\\_37/](https://www.hse.gov.uk/foi/internalops/hid_circs/permissioning/spc_perm_37/)

<sup>5</sup> (7388 fatalities in 2021)/(331M total population)

INDUSTRY	MAXIMUM INDIVIDUAL RISK (per year)	NEGLIGIBLE INDIVIDUAL RISK (per year)
Aircraft design (EASA)	-	-
ATM (EUROCONTROL)	-	-
Airports (UK)	10 <sup>-4</sup> (public)	10 <sup>-5</sup>
Road transport (EU MS)	-	-
Road transport (USA, Norway)	-	-
Road transport of DG (ACDS)	10 <sup>-3</sup> (workers), 10 <sup>-4</sup> (public)	10 <sup>-6</sup>
Road tunnels (Austria and others)	-	-
Rail transport (ERA)	Various FWSI per pass km	-
Rail transport (UK)	1.038 FWI per 10 <sup>8</sup> pass km	-
London Underground	10 <sup>-3</sup> (workers), 10 <sup>-4</sup> (public)	10 <sup>-6</sup>
Nuclear (ICRP)	10 <sup>-3</sup> (workers), 10 <sup>-4</sup> (public)	-
Onshore process (UK)	10 <sup>-3</sup> (workers), 10 <sup>-4</sup> (public)	10 <sup>-6</sup>
Onshore process (Netherlands)	10 <sup>-6</sup> (public LSIR)	-
Onshore process (Flanders)	10 <sup>-5</sup> (public LSIR)	10 <sup>-7</sup>
Onshore process (France)	-	-
Onshore process (HK)	10 <sup>-5</sup> (public LSIR)	-
Offshore oil & gas (UK)	10 <sup>-3</sup> (workers)	-
Healthcare	-	-

*Figure 2 – Individual risk criteria in different industries. Source: European Maritime Safety Agency<sup>6</sup>.*

- **Other exposed parties (thoughts for the long term)**

The advantage of measuring risk in the frame of reference of the exposed third party is that we can set an acceptable level of risk in line with the risk accepted by those parties.

For example, we could imagine setting an acceptable level of risk for astronauts in commercial space operations. In the hypothetical case that an astronaut accepts a risk of  $1 \times 10^{-4}$  per launch, it would be overly conservative to limit the HAPS risk to  $1 \times 10^{-9}$  per rocket launch. Instead, something closer to, but smaller than,  $1 \times 10^{-4}$  would be more appropriate.

Similarly, we could set an acceptable level of risk for another unmanned vehicle operator that is in line with the risk that this operator is accepting (e.g., Loon, a former unmanned free balloon operator, could tolerate safely losing a balloon every  $10^5$  hours).

While it is highly impractical to set an infinite list of acceptable levels of risks for each possible exposed party, new collaborative traffic management concepts (CTM) that are based on intent and information sharing make it possible for each operator to communicate its risk tolerance (which may depend on the location of the craft), so that the conflict identification resolution can adapt accordingly (see the Aeronautical Industries Association’s 2022 paper “Cooperative Operations In Higher Airspace A Proposal.”<sup>7</sup>)

In a collaborative traffic management framework where each operator shares its risk tolerance with intents, it becomes possible to apply larger safety margins to handle a conflict between a manned supersonic and an unmanned free balloon than for a conflict

<sup>6</sup> <https://www.emsa.europa.eu/publications/reports/download/3547/2419/23.html>

<sup>7</sup> <https://www.aia-aerospace.org/wp-content/uploads/AIA-Cooperative-Operations-in-Higher-Airspace-Proposal-April-2022-Final33.pdf>

involving two unmanned free balloons. Likewise, it becomes possible to apply additional safety buffers when a conflict is located over densely populated areas.

### ***Individual risk sharing between operators***

The risk to which one individual is exposed is the sum of the risk from all HAPS operating above them (which may be operated by different operators). This raises the question of risk sharing between operators. However, for the foreseeable future, only a handful of operators are expected to operate. Therefore, as a first approximation, we propose that the individual risk (which intrinsically embeds the number of HAPS) be specified per operator -- this avoids the complexity of risk sharing and attribution between HAPS operators. This can later be revised as needed should the number of operators scale dramatically. Interesting mechanisms can be considered:

- Cap-and-trade systems -- similar to that of CO2 emission regulations, airport slot allocation
- Auction systems and real-time bidding (e.g., Vickery-Clarke-Groves auction<sup>8</sup>) designed to assign items in a fairly and socially optimal manner

Current airspace operations already demonstrate concepts of sharing risk between operators, although it is usually not framed in terms of risk, but as capacity. Arrivals at airports are heavily constrained by airport capacity rates set to ensure separation minima are met and to accommodate reduced arrival rates during adverse weather. Complex ground-delay and airborne-delay systems are used to equitably distribute delay, which is a cost to the industry. Similarly, many arrival and departure slots are sold by airports to operators, recognizing the associated capacity and risk constraints. Here, commercial market forces dictate access to the scarce resource. A similar system would inevitably evolve as the density of HAPS increases.

## **Societal (Collective) Risk for HAPS**

The measure of individual risk discussed above does not account for population density or aircraft density. This is because the individual risk measures the risk to each exposed individual (or aircraft). The risk to one individual is not affected by the presence of other individuals.

Additional safety criteria may be needed to assess societal risk. Societal risk is defined as the relationship between frequency and the number of people affected by the harm in a given population from the realization of specified hazards. Societal risk limits have been defined in the transportation of dangerous goods<sup>9</sup>, by the UK HSE<sup>10</sup>, and a similar collective risk concept has been used by the FAA ALR approach for commercial space<sup>11</sup>. Societal risk criteria can be defined with F-N curves that specify the acceptable frequency (F) of an accident involving N or

---

<sup>8</sup> [https://en.wikipedia.org/wiki/Vickrey%E2%80%93Clarke%E2%80%93Groves\\_auction](https://en.wikipedia.org/wiki/Vickrey%E2%80%93Clarke%E2%80%93Groves_auction)

<sup>9</sup> <https://railroads.dot.gov/sites/fra.dot.gov/files/2020-02/Evaluation%20of%20Risk%20Acceptance%20Criteria.pdf>

<sup>10</sup> [https://www.hse.gov.uk/foi/internalops/hid\\_circs/permissioning/spc\\_perm\\_37/ - Tools-for-ALARP](https://www.hse.gov.uk/foi/internalops/hid_circs/permissioning/spc_perm_37/ - Tools-for-ALARP)

<sup>11</sup> [https://www.faa.gov/sites/aa.gov/files/space/additional\\_information/faq/SLR2\\_Final\\_Rule\\_450\\_2.pdf](https://www.faa.gov/sites/aa.gov/files/space/additional_information/faq/SLR2_Final_Rule_450_2.pdf)



more fatalities. Societal risk can also be measured using expectation value (the expected number of fatalities per year).

The ICAO SASP has traditionally not used F-N risk metrics, since their work has focused on larger regular public transport operations where the loss of an aircraft is always a significant number of fatalities. However, regulators often use F-N curves when assessing risk in regions where general aviation mixes with smaller operators; the probability of loss of a two-person recreational aircraft is treated differently from the potential loss of a 40-person regional operator.

One benefit of societal risk is that it incentivizes continuous safety improvement. It allows for a progressive approach to safety in which early R&D can benefit from low operational volumes and low-density operating areas to maintain acceptable levels of risk. As the systems mature and are proven more robust, it allows operators to gradually increase the density of population and aircraft overflown, and gradually increase operational volume (HAPS density), while monitoring that the collective risk always remains within the acceptable level.

Societal risk criteria can, however, be challenging to establish and do not scale with the societal value provided<sup>12</sup>. In particular, it can be difficult to translate societal risk from one industry to another, and it can be challenging to define an appropriate area over which societal risk should be accumulated.

We propose the following:

- To define societal risk metrics (and acceptable level of risk) as:
  - Risk to aircraft -- The (maximum) expected number of mid-air collisions for regular public transport<sup>13</sup> per year in a standard airspace grid
  - Risk to ground populations -- The (maximum) expected number of fatal accidents per year in a standard-size region
- To define a standard World grid over which the societal risk is aggregated for each operator (i.e. the societal risk is aggregated in each cell of the grid and kept below the acceptable level of risk such that no cell can exceed the acceptable level of risk). See grid example used by Loon for societal risk aggregation in figure 3.

To avoid overly restrictive regions, the grid must have cells of equal surface area or volume. The cell size must be sufficiently small to control societal risk at a local level (e.g. urban agglomeration) while keeping the risk aggregation of unrelated regions separate.

*Note 1:* risk acceptance is a function of the value gained from a service - As a HAPS operator expands its service to other distant regions, it provides additional service value and creates risk to a different population group. The societal risk should be aggregated separately, leveraging a sufficiently fine mesh grid, appropriately accounting for service value and risk with geographic expansion.

*Note 2:* separate grids, with different cell sizes, may be defined for air-risk and ground risk.

---

<sup>12</sup> <https://www.emsa.europa.eu/publications/reports/download/3547/2419/23.html>

<sup>13</sup> Regular public transport is intended to cover commercial flights and not small-sized general aviation aircraft.



- The acceptable levels of societal risk should be revised on a regular basis to remain in line with societal acceptance.

### **Societal risk sharing between operators**

Similar to individual risk sharing, we propose that, due to the limited number of HAPS operators expected, the ALR criteria be set for each operator (avoiding risk-sharing complexity). As the ETM/CTMS/ECHO concepts are developed further, mechanisms can be introduced to incentivize efficient airspace use, and efficient risk-budget use in high-demand regions. For example, cap-and-trade or bidding systems (even if non-monetary) could later be set to incentivize operators accessing high-demand airspaces to use the airspace and risk budget efficiently (lower-risk operations using less risk budget are cheaper to bid on).

## Summary of Proposed Acceptable Levels of Risk for HAPS

3rd-party-centric acceptable levels of risk	Individual Risk (set per operator)	Collective Risk (set per operator)
<b>Manned Aircraft</b>	$5 \times 10^{-9}$ mid-air collision per exposed aircraft flight hour	X mid-air collision per operator per airspace (std size grid) per year
<b>Ground Population</b>	$1 \times 10^{-4}$ - $1 \times 10^{-6}$ tolerable range probability of fatality per exposed person per year	Y fatality per region (std size grid) per operator per years

## Operator-Managed Societal Risk

An important benefit of using a societal risk criterion is that it can be easily computed (even in real-time) and operationally managed by operators in a way that is auditable by regulators. The following is needed:

- A standardized grid for summing the risk
- A standardized set of world population/aircraft traffic data is computed for that grid.
- **Platform-specific “risk factor” constants.** Indicators of the performance and risk of a specific design that is normalized per unit flight time and unit population density.

*The risk factors for a specific platform **could be obtained through an airworthiness/certification process**<sup>14</sup> or empirical flight data with mathematical*

<sup>14</sup> The purpose of certification process would evolve from a binary aircraft approval process, to that of certifying that a platform meets a specified risk rate.

*modeling or simulation.*

Example risk factors (imaginary numbers in the order of magnitude to that used by Loon):

- Air-risk factor --  $1.4 \times 10^{-13}$  mid-air collision per platform flight hour per aircraft density overflown (aircraft per square kilometer)
  - Ground-risk factor --  $3.8 \times 10^{-12}$  ground fatality per platform flight hour per population density overflown (people per square km)
- Timestamped historical fleet trajectories

The May 2023 ICAO SASP<sup>15</sup> meeting proposed a simple  $0.1 \times 0.1$  degree grid ( $\sim 6 \times 6 = 36$  NM<sup>2</sup>) with an allowance in the calculation for the actual area. This allows for simpler calculations than the use of hexagonal-like global grid systems.

With the above information, an operator can dynamically compute the societal risk across its entire fleet (e.g., via sum-product in a spreadsheet), and ensure that the risk budget is never exceeded in any grid cell.


Without needing operational approval for each new region, or every time the operator wishes to vary fleet density, the operator can flexibly adjust operational volumes, operational regions and overflown population/air traffic densities to ensure that the societal risk is never exceeded in any cells of the standard grid. Regulators can be confident that the total system risk is always maintained, and can audit operators by requesting historical fleet trajectory data. Operators could also share the computed societal risk such that regulators keep a real-time map of the total risk map.

Note: in the above framework, the meaning and purpose of the aircraft certification are changed from an aircraft approval binary process (approved to fly vs. not approved to fly) to a process validating that the “risk rate” used in collective risk computation and operational risk management reflects the reality. As such, the certification could certify an aircraft on a continuous risk scale:

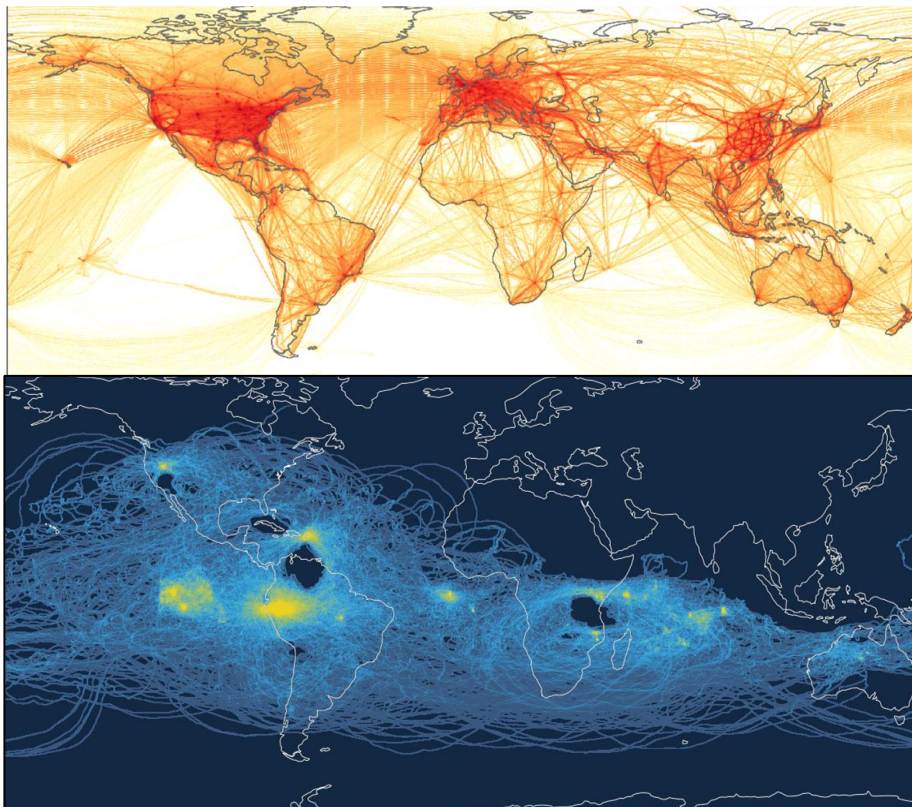
- A higher-risk rate for sustained fleet operation in low-risk regions with only occasional transits over denser regions
- A lower-risk rate for sustained, higher-density operations over higher population densities

---

<sup>15</sup> SI Barry, High Altitude Platform Systems: guidance material: Attachment on mathematical Modelling, ICAO Separation Airspace Safety Panel, Working Paper 4, Montreal, 19 May, 2023



**Figure 3** – Example grid and risk map computed by Loon to estimate the societal mid-air collision risk (in real-time) for the entire fleet. The color in each cell represents the number of years between expected mid-air collisions with aircraft (minimum cell value ~ 500,000 years). A similar grid was computed for ground risk.



**Figure 4** – (Top) Statistical aircraft density data used by Loon to compute Figure 3. (Bottom) Historical fleet position used by Loon to compute Figure 3.



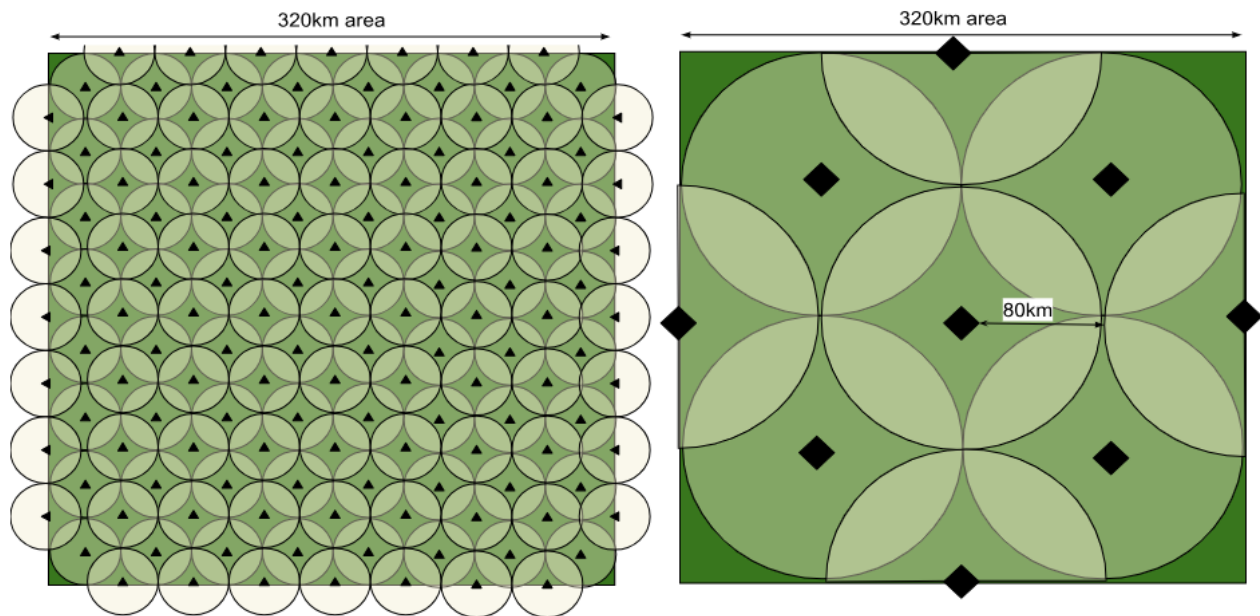
# Annex – Examples of Why Aviation Metrics Do Not Work

## Example #1 – Per-flight-hour risk metric is inadequate

Imagine a HAPS designer considering the following two platform options to provide connectivity over an area of 320km by 320km.

**Design A** is the smaller alternative, weighing 100kg, which can provide coverage over a radius of 20km. A platform of Design A has a likelihood of an unplanned descent of 1 in 100,000 flight hours. The likelihood that an unplanned descent results in a fatality on the ground can be estimated at 1 in 1,000 for the population density of the service area.

**Design B** is the larger platform alternative, weighing 1,000kg and capable of carrying a multi-beam payload that provides connectivity over a radius of 80km. **As a result of this larger coverage, Design B requires 16x fewer platforms to cover the service area than Design A.** A platform of Design B has a likelihood of an unplanned descent of 1 in 100,000 flight hours. Due to its larger size, the probability that an unplanned descent results in a fatality on the ground can be estimated at 1 in 100 for the population density of the service area.



**Design A** - 141 Smaller HAPS providing coverage over service area, each with connectivity radius of 20km

**Design B** - 9 Larger HAPS providing coverage over service area, each with connectivity radius of 80km

*Figure 5: Example of two possible design choices, one leveraging smaller HAPS operated at higher density, and another using larger HAPS with bigger coverage*

If we look at the risk per flight hour for the population density in the service area, we would conclude that Design A is 10x safer than Design B:



- **Design A**
  - $10^{-8}$  probability of ground fatality per HAPS flight hour (=  $1/100000 * 1/1000$ )
  - 1 ground fatality every 81 years<sup>16</sup>
- **Design B**
  - $10^{-7}$  probability of ground fatality per flight hour (=  $1/100000 * 1/100$ )
  - 1 ground fatality every 127 years<sup>17</sup>

Design B is however a safer choice when looking at the operation holistically, despite having a risk per flight hour 10x greater than Design A.

This example illustrates how a TLS defined on a per-flight-hour basis (or a Type Certification Process that focuses on a per-flight-hour basis) could miss the big picture and incentivize HAPS manufacturers to opt for a design that is less safe than a disqualified alternative.

Note: The monitoring method proposed at the May 2023 SASP meeting only considered the risk to aircraft flying underneath a HAPS. Here, the number of aircraft points (each a five-second sample) in a month (or similar long period), is recorded in each  $0.1 \times 0.1$  degree grid; hence the risk over time for any operation can be calculated as proportional to the number of HAPS points times the number of aircraft points. The real-time accumulation of risk allows owners to modify HAPS operations to balance total risk against mission goals. The total risk can be measured against the number of aircraft hours in the overall region ( $< 10^{-8}$  collisions per flight hour), some reasonable measure of years between collisions or some measure of collisions per flight. The SASP work recognized that what constitutes the “region” and hence “years between collisions” is not absolute, and will need to be individually considered by the regulator.

## Example #2 - Per-mission risk metric is inadequate

Imagine a HAPS designer who is considering the following two system designs for operating a single HAPS continuously:

**Design A** uses advanced materials that make it capable of staying aloft for an entire year. It has a probability of an unplanned descent of 1 in 1,000 missions. Each of these unplanned descents has a probability of 1 in 1,000 to generate a mid-air collision with manned traffic operating below.

**Design B** uses different materials such that the platform can only remain airborne for 1 month at a time. To maintain continuous connectivity service during the year, two platforms are used and cycled each month such that there is always one airborne platform while the other one is in maintenance. The probability of unplanned descent for Design B is 1 in 5,000 missions, and each of these unplanned descents has a probability of 1 in 1,000 to generate a collision with manned traffic operating below.

On a per-mission basis, Design A has 5x more risk of mid-air collision than Design B. However, when looking at the system holistically, Design A is the safer choice.

---

<sup>16</sup>  $1/(365\text{days} * 24\text{hours} * 141 \text{ platforms} * 10^{-8} \text{ fatality/hour}) = 81 \text{ years between fatalities}$

<sup>17</sup>  $1/(365\text{days} * 24\text{hours} * 9 \text{ platforms} * 10^{-7} \text{ fatality/hour}) = 127 \text{ years between fatalities}$

- **Design A**
  - $1 \times 10^{-6}$  mid-air collision per mission (=  $1/1000 \times 1/1000$ )
  - 1 mid-air collision expected on average every 1 million years<sup>18</sup>
- **Design B**
  - $2 \times 10^{-7}$  mid-air collision per mission (=  $1/5000 \times 1/1000$ )
  - 1 mid-air collision expected on average every 416 thousand years<sup>19</sup>

A HAPS manufacturer may therefore elect Design B (or be constrained to do so if target safety levels were specified on a per-mission basis).

We can see from this example how a TLS defined on a per-mission basis (such as proposed by EASA's draft) can be misaligned with true safety objectives for HAPS, and could incentivize HAPS manufacturers to opt for designs that optimize for that metric rather than optimizing for overall safety.

## Join Us in Our Work

All companies interested in the HAPS ecosystem are encouraged to become [HAPS Alliance members](#). Alliance membership levels – Principal, General, and Supporter – are open to organizations in any industry sector. Principal and General Members have the opportunity to become involved in various membership initiatives, including [working groups](#), [member-only meetings](#), and collaboration with other HAPS Alliance members to work on technology components and use cases for enabling a smarter world. View all of the [member benefits](#) and see why you should [join](#) today.

## About the HAPS Alliance

The HAPS Alliance is an industry association of High-Altitude Platform Station (HAPS) industry leaders that include telecommunications, technology, aviation and aerospace companies, as well as public and educational institutions. United by a vision to address diverse social issues and create new value through the utilization of high-altitude vehicles in the stratosphere, the Alliance is working to accelerate the development and commercial adoption of HAPS technology by promoting and building industry-wide standards, interoperability guidelines and regulatory policies in both the telecommunication and aviation industries. For more information, please visit <https://hapsalliance.org/>.

---

<sup>18</sup>  $1/(10^{-6} \text{ collision per mission} \times 1 \text{ mission per year}) = 1\text{M years between collisions}$

<sup>19</sup>  $1/(10^{-7} \text{ collision per mission} \times 12 \text{ missions per year}) = 416\text{k years between collisions}$