

Assurance and accreditation of autonomous systems in Australia: Current State Body of Knowledge

Current Australian frameworks for assurance and accreditation of autonomous systems in the Air Domain

All inquiries can be addressed to info@tasdcrc.com.au, Trusted Autonomous Systems, Queensland, Australia.

Disclaimer

The Trusted Autonomous Systems Defence Cooperative Research Centre accepts no liability for the accuracy of the information nor its use or the reliance placed on it.



Table of Contents

1	Air domain: Current Australian framework for assurance and accreditation of autonomous systems	4
2	An introduction into Australian Aviation (from a legislative perspective)...	5
2.1	Introduction into general state of aviation legislation.....	5
2.2	International obligations under the Chicago Convention.....	6
2.3	The high level structure of Australian Aviation Legislation.....	6
2.4	CASA's Regulatory Philosophy	8
2.5	How the Australian legislation breaks down aviation operations	8
2.6	Aircraft engaging in international air navigation.....	12
2.7	The constituent responsibilities and requirements that allow for safe operations to emerge.....	12
3	Detailing the regulatory requirements and then deriving the high-level safety objectives for Autonomous Aircraft and RPAS.....	15
3.1	The Civil Aviation Act Requirements and how the flow on effects to uncrewed aircraft system requirements	15
3.2	Part 101 of CASR and the General Structure of UAS regulation.....	16
4	Understanding the requirements for each constituent responsibility and requirement that create safety.....	20
4.1	A discussion on Operational Procedures and Flight Crew regulatory requirements for complex UAS operations	20
4.2	A discussion on the airworthiness regulatory requirements for complex UAS operations	21
4.3	A discussion on Airspace Integration regulatory requirements for complex UAS operations.....	26
5	Aircraft operating for Defence (State Aircraft)	30
5.1	The DASR, Applicability, Structure and Decomposition of Aviation Activities	30
6	Comparative studies – other countries integration of UAS	36
6.1	European Union – European Union Aviation Safety Agency	36
6.2	United States – Federal Aviation Administration	38
7	Current methods for assurance of highly automated/autonomous UAS (JARUS SORA).....	41
7.1	Safety Risk Management	42
7.2	Systems Safety Engineering in Aviation	44
7.3	Measuring risk for unpersonned operations	45
7.4	The Aviation Safety Continuum.....	48
7.5	Acceptable levels of safety for aviation.....	49
7.6	Acceptable levels of safety for third parties on the ground	49
7.7	Acceptable levels of safety for third parties in the air	49
7.8	Air Risk and Ground Risk as a function of the Concept of Operations	51
7.9	Bow-tie analysis to visualise hazardous events for UAS	55
7.10	The system safety process as a model for safety assurance of technical systems	60
7.11	The SORA Risk Assessment, the SAIL and the Aviation Safety Continuum	61

8	Assurance issues for highly automated and autonomous UAS	65
8.1	Incorporating analysis of autonomy into the SORA and system safety process.....	66
9	References	74

1 Air domain: Current Australian framework for assurance and accreditation of autonomous systems

This document provides a comprehensive discussion on the current state of assurance and accreditation of autonomous systems in the Air Domain.

The intent is to provide an end-to-end discussion, starting with the reasons for aviation legislation, and concluding with a discussion on a currently accepted safety assurance process and the gaps, issues and opportunities that could be implemented to better integrate highly automated and autonomous aircraft into Australia airspace.

This document provides a comprehensive discussion on the current state of assurance and accreditation of autonomous systems in the Air Domain.

The intent is to provide an end-to-end discussion, starting with the reasons for aviation legislation, and concluding with a discussion on a currently accepted safety assurance process and the gaps, issues and opportunities that could be implemented to better integrate highly automated and autonomous aircraft into Australia airspace.

- **2.0** introduces the Australian aviation legislative context, explaining the international pressures that led to internationally harmonised and agreed policies for international aviation and how this results in the implementation of the Civil Aviation Act 1988, Australia's main piece of legislation that regulates the UAS industry (from a safety perspective). The regulatory framework and context across all aviation activities (and how it applies to UAS) is also detailed.
- **3.0** presents a brief discussion on how the current assurance requirements for UAS boil down to demonstrating that an operation does not have an adverse effect on the safety of air navigation.
- **4.0** then follows the legislative flow from high level legal requirements for aviation into UAS specific accreditation requirements for the different constituent roles and responsibilities that create safe aviation activities possible.
- **5.0** introduces the rules, regulations, and the regulator (DASA) for Australian State Aircraft. This section particularly highlights the differences in approaches between DASA and CASA when regulating aviation activities.
- **6.0** provides a comparative study between the two largest international civil regulators of aviation activities, the United States Federal Aviation Administration (FAA) and the European Union Aviation Safety Agency (EASA). This section will focus on how UAS are regulated by these entities and the key enabling work being undertaken by these two regulatory giants to better integrate UAS operations into the national airspace systems.
- **7.0** details how assurance can be undertaken for UAS, providing models for assessing ground and air risk posed to third parties, alongside a casual bow-tie model that can be used to understand how risk controls and mitigations affect the likelihood of a hazard to these third parties. The SORA process is introduced as an end-to-end, risk proportionate safety assurance process that brings together the elements discussed earlier.
- **8.0** identifies the gaps, challenges and opportunities that need to be addressed in current processes in order to truly integrate autonomous systems into the Australian aviation industry.

2 An introduction into Australian Aviation (from a legislative perspective)

This section is intended to provide a comprehensive overview of the current Australian aviation legislative context, the scope of aviation activities and the associated components that need to work together properly to ensure safety.

2.1 Introduction into general state of aviation legislation

In the early 20th century, rigid airships were becoming larger and more capable, and heavier than air flight was beginning to emerge as a viable way to take to the skies. During the first decade of this century, many European states began to question how they would deal with transnational operations of aircraft. With land and maritime transport, there were natural or man-made barriers that could be put in place to restrict or allow operations of vehicles and vessels between nations, but no barriers existed in the air.

The first international conference on navigation was held in Paris in 1910 to attempt to come to some international agreement on how to effectively deal with international aviation. By 1919 the first agreed International Convention for Air Navigation (ICAN) was developed, containing 43 articles and guidance for registration, certificates of airworthiness, pilot licencing etc. that states should follow to ensure safe and open international travel.

In 1944, towards the end of the second world war, it was clear to all nations that international aviation would be accelerated by the aviation technology developed during the war. In order to, “preserve friendship and understanding... [and] avoid friction and promote cooperation between nations and peoples upon which the peace of the world depends”¹, the convention on international civil aviation was signed in Chicago. The convention was ratified in 1947 and the International Civil Aviation Organisation (ICAO) was established.

Signatories to the convention have, “agreed on certain principles and arrangements in order that international civil aviation may be developed in a safe and orderly manner and that international air transport services may be established on the basis of equality of opportunity and operated soundly and economically”.² ICAO facilitates the intent and mission of the convention.

It needs to be noted that ICAO is not a regulator and cannot supersede a signatory’s national regulator. Rather ICAO exists to ensure harmonisation and cooperation towards safe outcomes for international aviation. As a signatory to the Chicago Convention, Australia has made legislation throughout the years that ratifies this international obligation (the Air Navigation Act, Civil Aviation Act) and places obligations on the operation of Australian aircraft and operations within Australian Airspace.

¹ ICAO, “Convention of International Civil Aviation.”

² ICAO.

2.2 International obligations under the Chicago Convention

The Civil Aviation Safety Authority (CASA) is Australia's Civil Aviation Authority. CASA was created under the *Civil Aviation Act 1988* (the Act) to ensure the safety of civil air operations within Australia and of Australian aircraft overseas³ in compliance with Australia's international obligations under the Chicago Convention on International Civil Aviation, signed by Australia in 1944. Effectively the Act gives CASA authority to regulate Australian aircraft, which includes **any non-foreign registered or state aircraft** flown in Australian territory and overseas, and the authority to regulate the operation of civil, non-state aircraft within Australian territory.⁴ An uncrewed or unpersonned aircraft is an aircraft for the purposes of the Act.⁵

A critical concept within the Chicago Convention is the delineation between responsibilities for aircraft operations and the aircraft themselves, dependent on whether the aircraft is operating within a State (defined as within 12NM⁶ of the state's territory) or over the high seas.

The Chicago Convention, signed by 191 countries⁷, requires that signatories ensure:

- When operating within the 12NM of a State, any aircraft must comply with the rules of that State; and
- When operating over the high seas (outside 12NM from a countries land), any aircraft must comply with the Rules of the Air over the high seas (Article 12 to the Convention), which are being updated to incorporate remotely piloted aircraft systems into Part IV to Annex 6 of the convention; and
- Aircraft have the nationality of the State in which they are registered (Article 17 to the Convention). Under Articles 31 and 33 of the Convention, every state must have issued a valid Certificate of Airworthiness for any aircraft from that state engaging in international operations and must accept any Certificate of Airworthiness issued by a state if it meets the minimum requirements of the Convention (which is covered in Annex 8 to the Convention). Effectively this means that a State is responsible for the airworthiness of aircraft registered to that state.

2.3 The high level structure of Australian Aviation Legislation

Australia's aviation legislation is comprised of three main tiers, as shown in Figure 1. The Act is the primary tier of legislation and while it is not the only aspect that CASA must consider, it is the most important one in this context. The international obligations mentioned in **0** drive the direction of Australian aviation legislation, which is captured in ss. 9(1)(a) and ss. 9(1)(b) of the Act:

9(1) CASA has the function of conducting the safety regulation of the following, in accordance with this Act and the regulations:

³ ss. 9(1)(a) and (b) of the Civil Aviation Act 1988

⁴ s. 3 of the Act, however for the majority of Regulations, State Aircraft are exempt as per s. 4 of the Act

⁵ As the definition in s. 3 of the Act states, an aircraft is any machine or craft, "that can derive support in the atmosphere from the reactions of the air, other than the reactions of the air against the earth's surface"


⁶ United Nations, United Nations Convention on the Law of the Seas.

⁷ ICAO, "Convention of International Civil Aviation."

- a. civil air operations in Australian territory
- b. the operation of Australian aircraft outside Australian territory

Along with many other subsequent functions to support this, further mentioned within s. 9 of the Act. The Act enables two secondary tier regulations, *the Civil Aviation Regulations 1988* (CAR) and *the Civil Aviation Safety Regulations 1998* (CASR). The CASR is slowly replacing the older CAR regulations, however at this point in time, part of CAR are still extant and applicable to certain aviation activities.

Figure 1 - Legislative Tiers in Aviation



Underneath the secondary legislation there are three important pieces of tertiary legislation, Civil Aviation Orders, or CAOs (enabled under CAR), Manual of Standards or MOS (enabled under CASR) and directions (enabled under Subpart 11.G of CASR).

Each subsequent tier of legislation is enabled by the former, which will become an important consideration when determining the overarching safety requirements for risk-based operations.⁸

⁸ Any contradiction between a parent (higher tier) and child (lower tier) regulation is resolved by the parent legislation taking precedence

2.4 CASA's Regulatory Philosophy

Although legislation is critical, there are additional policies and philosophies that need to be taken into consideration whenever interpreting the regulatory requirements. Key considerations for this document are the following points from CASA's Regulatory Philosophy:⁹

Principle 3. CASA takes risk-based approaches to regulatory action and decision-making

Principle 4. CASA performs its functions consistently with Australia's international obligations

Principle 7. CASA fairly balances the need for consistency with the need for flexibility


These three points effectively act the lenses through which CASA interprets legislation and assists in the determination of appropriate means of compliance.

2.5 How the Australian legislation breaks down aviation operations

Drawing on the previous section's explanation of the reasons for the Act, we can begin to look at how the Act and subsequent regulations delineate different aviation activities. This will help establish an appreciation for the scope of possible aviation operations and identify where this body of knowledge will focus.

2.5.1 First Major Delineations – Whether the Operation uses a Rocket, a State Aircraft, a Crewed Civilian Aircraft or an Uncrewed Civilian Aircraft

Figure 2 - Aviation Activities in Australian Airspace (Decision Tree 1)



To begin, we look at the following delineations, graphically represented above:

⁹ CASA, "Our Regulatory Philosophy."

1. If the activity utilises an aircraft (or not)

This is a technical delineation, as the Act imposes significant requirements on **aircraft** which are defined as “any machine or craft that can derive support in the atmosphere from the reactions of the air, other than the reactions of the air against the earth’s surface”. Critically, a rocket does not fall under this definition, and whilst subject to Part 101.H of CASR, does not need to meet significant parts of the regulations due to this definition. Rocketry operations will not be covered in this analysis.

2. If the activity is undertaken by a State aircraft or a civil aircraft

This is a critical regulatory oversight definition, as a State aircraft is not usually regulated by CASA. Foreign State aircraft (an aircraft used in the military, customs or police services of a foreign country) are regulated by their respective state (however they must follow Australian operational rules) and will not be catered for in this report. Australian Defence Force aircraft are also State aircraft and are usually regulated by the Defence Aviation Safety Authority (DASA), under the Defence Aviation Safety Regulations (DASR) which is discussed in 0.

All other aircraft are civil aircraft and are regulated by CASA in significant capacity. This includes Australian aircraft and foreign civil aircraft.

3. If the activity is a crewed or uncrewed aviation activity

Given that the aircraft is a civil aircraft, the next major delineation is whether the operation is crewed or uncrewed. This split drives whether the operation is covered by the traditional, conventionally piloted aircraft regulations, or by uncrewed aircraft regulations under Part 101 of CASR.


To be clear, the current aviation regulations uses the term “Unmanned Aircraft”, which is termed here as “uncrewed”. Unmanned Aircraft Systems (UAS) are aircraft¹⁰ that are intended to be operated with no pilot onboard.¹¹ “Uncrewed” is a more useful term, as the term “unmanned” implies no people on board, which does not align with the CASA and ICAO definition (i.e. an UAS can be an aircraft that carry passengers onboard, but no crew). Additionally, the term “uncrewed” is not gendered.

Although this report will not focus on crewed aviation activities, highly automated/autonomous systems will be pervasive throughout these operations currently, presently, and into the future. However, as these complex, highly-automated systems are (generally) covered by ICAO regulations (or are not at all automated – see Part 105 and 132 of CASR) and by stringent certification requirements. For completeness, the different crewed activities are shown in decision tree 3 below.

¹⁰ As defined in s. 3 of the Act, an aircraft “means any machine or craft that can derive support in the atmosphere from the reactions of the air, other than the reactions of the air against the earth’s surface”

¹¹ ICAO, “Unmanned Aircraft Systems (UAS).”


Figure 3 - Crewed Aviation Activities in Australian Airspace (Decision Tree 3)



2.5.2 Uncrewed Aircraft Activity Delineations

Under the umbrella term of Uncrewed Aircraft Operations, we can further delineate the operations based on 4 criteria, shown in Decision Tree 2:

Figure 4 - Decision Tree 2 (DT.2) Uncrewed Aircraft System Delineations



1. Whether the aircraft is an unmanned free balloon or an uncrewed aircraft operation

This delineation is important because unmanned **free** balloons do not have any propulsion systems (they may still have systems to control altitude), whilst aircraft have substantially more control over their position, velocity and orientation in space and time. This analysis will not consider unmanned free balloons.

2. Whether the aircraft is autonomous, a model remotely piloted aircraft, an excluded RPA or an “included RPA”

Under the umbrella term of Uncrewed Aircraft System, there are the following categories and subcategories to consider:

- Remotely Piloted Aircraft Systems (RPAS)

i.e. those unmanned aircraft systems that are affected by the pilot/crew from offboard the aircraft. The aircraft itself is defined as the RPA, whilst the RPA and all supporting infrastructure (i.e. remote pilot station, communications infrastructure) is defined as the RPAS.

- Model Aircraft RPAS
- Excluded Category RPAS Operations
- “Included” Category RPAS Operations

Part 101 of CASR specifies the rules and requirements for unmanned aircraft and rockets. These requirements take the form of restricting the scope of operations unless certain approvals are obtained. Model aircraft and excluded category operations are, by definition, extremely constrained operations, which restricts the capability of systems

For the purposes of this analysis, we will be focusing on complex operations that require a remotely piloted operating certificate (ReOC) and additional approvals under Part 101 (termed “included” RPAS operations). The excluded and model aircraft categories will not be discussed here.

- Autonomous Aircraft

i.e. an unmanned aircraft that, “does not allow pilot intervention throughout all phases of the flight of the aircraft.”¹²

The definition of autonomous aircraft in CASA parlance is on the extreme end of what most people consider autonomous. The analysis in this report **will** take into account autonomous aircraft.

It is important to note that for the purposes of the Body of Knowledge, the term “autonomous” does not align with this legal definition. This means that an autonomous

¹² Subregulation 101.097(1)(a) of CASR

aircraft in the Body of Knowledge could be an RPA **or** an autonomous aircraft under the definitions in Australian legislation.

2.6 Aircraft engaging in international air navigation

The Chicago Convention requires that signatories ensure:

- When operating within 12NM of a State's land, any aircraft must comply with the rules of that State; and
- When operating over the high seas (outside 12NM from a State's land) each aircraft must comply with the Rules of the Air over the high seas (Article 12 to the Convention) convention.

This effectively means that without a specific approval to operate in a manner contrary to the international rules or rules of an overflowed state, any aircraft operation must comply with the standards for international aviation, which are covered particularly in Annexes 2, 6, and 8 to the Convention. These annexes are currently being updated to include operational, training and airworthiness requirements for international RPAS operation. It is important to note that at this point in time ICAO has specifically excluded autonomous aircraft¹³ from being covered under the international standards.

Until these updates are incorporated, it is likely that one-off approvals for highly automated aircraft organised between states and operators will be the norm. In the short term it appears that international autonomous aircraft operations will not be permitted.

2.7 The constituent responsibilities and requirements that allow for safe operations to emerge

Whilst it may seem that delineating aviation activities is easy from the previous section, there is more to the story. In each activity there is a complex interrelationship between people, systems, organisations, and processes that need to be regulated to ensure appropriate levels of safety for each of these activities. The CASR Regulatory Structure Webpage has the following text that tries to capture the elements that make a flight safe:¹⁴

Safe flights use **registered** and **certificated aircraft**, **maintained** by a **qualified person**, **flown** by a **qualified person** in **airspace** controlled by a **qualified person**.

Where the different colours indicate the different CASR Parts which create responsibilities and requirements for each part. This sentence is heavily focused on

¹³ 2.2 of ICAO Circular 328 (AN/190)

¹⁴ CASA, "CASA's Regulatory Framework."

traditional, conventionally piloted aviation in controlled airspace (no qualified person controls Class G Airspace).

A more broad and useful statement would be:

Safe Operations use **Registered** (when needed) and **Safe-for-Flight Aircraft**, **Maintained Following Appropriate Procedures** and Released for flight by a **Qualified Person**, **Flown Following Appropriate Procedures**, overseen by a **Qualified Person** in **Adequately Protected Airspace**.

What is important to note here that although the main delineation between requirements for different activities is stipulated through the operational, these are supported by a series of regulatory responsibilities and requirements regarding the **Airworthiness**, **Operations** and **Airspace Integration** of these operations, and are further elaborated below in **Table 7**.

These will be referred back to when discussing the current regulatory requirements for RPAS operations in each of the Airworthiness, Operations, and Airspace Integration aspects of safe aviation activities.

Table 1 - Constituent parts of a safe aviation activity

	Term used	Description
Airworthiness (Initial and Continuing)	Safe Operations use	Operations that meet a level of safety that is acceptable to the regulator and society at large.
	Registered Aircraft	An aircraft that has been placed on the Australian Aircraft Register if required to do so under Part 47 of CASR (some aircraft are exempt).
	and Safe For Flight Aircraft,	An aircraft that is safe for flight is one that meets a given level of safety for the type of operation. At higher levels of risk, this will include that a particular aircraft undertaking flight <u>conforms</u> to an <u>approved design</u> and is in a <u>condition safe for flight</u> (defined as airworthy). ¹⁵ At low levels of risk, an aircraft may not need to be fully airworthy (i.e. there may not need to be an approved design), but should still be in a condition safe for flight.
	Maintained Following Appropriate Procedures	In this use of the phrase, maintenance is the Continuing Airworthiness process, which includes both; when to undertake maintenance (scheduling, periodical inspections etc.) and how maintenance is undertaken on an aircraft. This process is needed to ensure that while an aircraft is in operation, it continues to be in a condition for safe flight (and in the context of airworthy aircraft, that the aircraft continues to comply with its approved design).
	and released for flight by a Qualified Person ,	A key safety responsibility (from a liability perspective) is a responsible person who “releases” the aircraft from maintenance. This person “certifies” that the aircraft has correctly completed all necessary maintenance and is now in a safe state for flight operations.
Operations	Flown Following Appropriate Procedures,	Aircraft need to be operated by appropriate procedures commensurate to the operational risk. An operator of a small drone for photography does not need to follow similar, complex procedures as a pilot of a large transport aircraft, but should achieve similar, safe outcomes
	overseen by a Qualified Person	Another key safety responsibility is the person who is responsible for the safety of the flight (i.e. the pilot in command or the remote pilot in command). Even in a fully autonomous operation, there will still be a person that takes responsibility for the safety of the operation, as this liability must be held by a person.
Airspace Integration	in Adequately Protected Airspace	Different classes of airspace have different equipage and procedure requirements. These requirements exist to protect all airspace users to a level that is accepted by society. In some cases, this will have a person responsible for controlling the airspace, and in others there may be no one responsible for controlling the airspace and responsibility for safety lies with the individual aircraft operating in that area.

¹⁵ See Regulation 42.015(2) of CASR 1998

3 Detailing the regulatory requirements and then deriving the high-level safety objectives for Autonomous Aircraft and RPAS

This section is intended to explain how the set of safety regulations for highly automated RPAS or autonomous aircraft are driven from the Civil Aviation Act and show their underlying outcome-based requirements.

3.1 The Civil Aviation Act Requirements and how the flow on effects to uncrewed aircraft system requirements

The Act and its effect on the safety of air navigation can best be contextualised by focusing on the following sections of the Act (this list is not exhaustive):

- The overarching safety requirements within the Act:
 - ss. 9(1) (CASA's functions),
- Specific requirements to be met within the Act
 - ss. 20AA(1) (Flying unregistered aircraft)
 - ss. 20AA(3) (Flying without a certificate of airworthiness)
 - ss. 20AA(4) (Flying without satisfying safety requirements)
 - ss. 20AB(1) (Undertaking critical without licence)
 - ss. 20AB(2) (Carrying out maintenance without authorisation)
- s. 98 (Regulations) of the Act.

ss. 9(1) of the Act states that “CASA has the function of conducting **the safety regulation** of... civil air operations in Australian territory... [and] ...the operation of Australian aircraft outside Australian territory”.

This implies the elements of conducting civil air operations and the operation of Australian aircraft, which are further decomposed in Part III of the Act (Regulation of civil aviation). In particular s. 20AA and s. 20AB which requires certain authorisations to be granted (or specifically be written that an authorisation is not required) in order to conduct civil air operations.

s. 98 grants the Governor-General the power to create regulations which further refine and define requirements for:

- (a) prescribing matters required or permitted by this Act to be prescribed;
- (b) prescribing matters necessary or convenient to be prescribed for carrying out or giving effect to this Act;
- (c) for the purpose of carrying out and giving effect to the provisions of the Chicago Convention relating to safety;
- (d) in relation to safety of air navigation within a Territory or to or from a Territory;
- (e) in relation to safety of air navigation, being regulations with respect to trade and commerce with other countries and among the States; and
- (f) in relation to safety of air navigation, being regulations with respect to any other matter with respect to which the Parliament has power to make laws.

These sections of the Act provide a basis for which CASR are applied and interpreted.

3.2 Part 101 of CASR and the General Structure of UAS regulation

To provide some logic to Part 101, we can break down the approvals into authorisations required for the purposes of Part III of the Act, as well as generic authorisation requirements (such as the ReOC and licencing for pilots) that must be acquired in order to undertake complex operations. These approvals that allow for an applicant to meet the requirements are delineated by whether they apply for Large RPA or non-Large RPA (i.e. very small, small and medium RPA).

These general approvals allow for an operator to operate a micro, very small, small, or medium (i.e. not a large RPA) RPA (i.e. not an autonomous aircraft) within the following restrictions:

- Operations must be below 400ft above ground level (AGL)
- Operations must be generally within Visual Line of Sight (VLOS)
- Operations must occur in the following day/weather conditions:
 - Within Visual Meteorological Conditions (VMC)
 - During the day
 - Not in or into cloud
- Operations must be conducted more than 3NM from the movement area of an aerodrome
 - For controlled aerodromes and non-controlled aerodromes, Chapters 4 and 9 of the Part 101 Manual of Standards provides more detail on the specifics of operating within 3NM of a movement area of an aerodrome for controlled and non-controlled aerodromes respectively
- Operations must be conducted more than 30m from people
 - This requirement depends on whether the operation is undertaken using an airship or the person has given consent for closer operation (greater than 15m)

- Operations must not be conducted over populous areas

Figure 5 - Linking the Act and CASR


the Civil Aviation Act 1988 Requirements to be met		Approval Regulations for the Act (very small, small and medium RPA)	Approval Regulations for the Act (pertaining to large RPA)
ss.20AA(1) Flying unregistered aircraft etc.	→	CASR 47.096A Certain RPA and model aircraft - registration requirements	CASR 47.015 Requirements for aircraft to be registered
ss.20AA(3) Flying without a certificate of airworthiness etc.	→	CASR 101.046 (Subpart 101.AB) flying without a certificate of airworthiness	CASR 101.255 (Subpart 101.F) Large RPA - requirement for certificate
ss.20AA(4) Flying without satisfying safety requirements	→	CASR 101.340 (Subpart 101.F) Conditions on certification	CASR 101.260 (Subpart 101.F) Maintenance of large RPA
ss.20AB(1) Undertaking essential duty without authorisation	→	CASR 101.048 (Subpart 101.AB) Performance of duties during flight time - remotely piloted aircraft	
ss.20AB(2) Carrying out maintenance without authorisation	→		CASR 101.265 (Subpart 101.F) Large RPA - persons permitted to carry out maintenance
CASR 1998 Requirements to be met		Approval Regulations for the Act (very small, small and medium RPA)	Approval Regulations for the Act (pertaining to large RPA)
CASR 101.252 (Subpart 101.F) Certain RPA - Requirement for remote pilot licence	→	CASR 101.295 (Subpart 101.F) Eligibility for remote pilot licence	
CASR 101.270 (Subpart 101.F) Certain RPA - Requirement for RPA operator's certificate	→	CASR 101.335 (Subpart 101.F) Eligibility for certification as RPA operator	

Further specific approvals are required to operate in higher risk environments. The approvals needed are shown below (note that these approval regulations are also the same regulations that prohibit operating in the specified context without an approval). There are effectively three different approval regulations that are utilised in Part 101:

- Approvals for Part 101 (101.029 of CASR)
- Area Approvals for Part 101 (101.030 of CASR)
- Approvals subject to Regulation 11.055 of CASR

The legend in the below figure delineates the type of risk that is affected by the approval (generically increases risk, affects ground risk or air risk):

Figure 6 - Specific approvals for Part 101



All of these approvals are subject to subregulation 11.055(1A) of CASR. This subregulation provides for the overarching requirements to be met in order for CASA to meet its obligations under the Act. Key requirements to be met in order for CASA to issue an approval under these regulations are highlighted below:

Table 2 - Subregulations in the CASR pertinent to the approval of UAS operations

Subregulation	Text
101.030(3)	In considering whether to approve an area for any of those purposes, CASA must take into account the likely effect on the safety of air navigation of the operation of unmanned aircraft in, or the launching of rockets in or over, the area.
101.030(6)	CASA may impose conditions on the approval in the interests of the safety of air navigation .
101.030(8)	CASA may revoke the approval of an area, or change the conditions that apply to such an approval, in the interests of the safety of air navigation , but must publish details of any revocation or change in NOTAM or on an aeronautical chart.
11.055(1A)(e)	granting the authorisation would not be likely to have an adverse effect on the safety of air navigation .
11.056(1)(a)	CASA may grant an authorisation subject to any condition that CASA is satisfied is necessary: for an authorisation, other than an authorisation to which subregulation 11.055(1B) applies or an experimental certificate—in the interests of the safety of air navigation

Based on these regulations, any approval for unmanned aircraft must ensure that there is not an adverse effect on the safety of air navigation (note that this requirement must be met alongside any prescriptive requirement in the regulations and that this discussion does not constitute the only requirement to be met).

Due to the relatively outcome-based style of Subpart 101.C, and 101.F, the requirement to ensure the safety of air navigation, becomes a critical. Of course, this leads to the next logical question, “what is the accepted baseline for the ***safety of air navigation?***”.

4 Understanding the requirements for each constituent responsibility and requirement that create safety

This section moves the discussion towards how the constituent responsibilities and requirements that make up the foundations for safe UAS operations.

Safe Operations use **Registered** (when needed) and **Safe-for-Flight Aircraft, Maintained Following Appropriate Procedures** and Released for flight by a **Qualified Person, Flown Following Appropriate Procedures**, overseen by a **Qualified Person** in **Adequately Protected Airspace**.

The following sections will delve deeper into the specific responsibilities and requirements for UAS under Part 101 and highlight some key opportunities for change that will further enable highly automated and autonomous operations.

4.1 A discussion on Operational Procedures and Flight Crew regulatory requirements for complex UAS operations

Safe Operations [are] use **Registered** (when needed) and **Safe-for-Flight Aircraft, Maintained Following Appropriate Procedures** and Released for flight by a **Qualified Person, Flown Following Appropriate Procedures**, overseen by a **Qualified Person** in **Adequately Protected Airspace**.

Operators of included RPAS operations (which include approvals for autonomous aircrafts) are required to obtain a ReOC and have:

- Appropriate operating procedures for the type of operation being undertaken;
- Appropriately qualified Chief Remote Pilot (CRP); and
- A maintenance controller undertaking duties described in **4.2.3**.

The scope of what constitutes appropriate procedures, and an appropriately qualified Chief Remote Pilot depends on the complexity and type of the operation. The ReOC forms the basis of any approval for operating UAS operations and additional approvals are given to an operator in addition to their ReOC.

Additionally, any person operating the RPA must have a Remotely Piloted Aircraft Licence (RePL) for operating the RPA(s) involved in the specific operation. The RePL is usually delineated by the category (i.e. whether the licence is for an aeroplane, rotorcraft (single or multi rotor) or powered lift) and the classification or weight class of the RPA (i.e. micro, very small, small or medium RPA).

Of particular note is if the operation utilises a large RPA,¹⁶ which requires the operator to obtain an approval under Regulation 101.275 of CASR 1998, which may require additional training and pilot competencies to operate these large RPA safely.

Readers are encouraged to utilise the resources for [operators](#) and [remote pilots](#) on the CASA website for additional information and guidance.

4.2 A discussion on the airworthiness regulatory requirements for complex UAS operations

Safe Operations use Registered (when needed) and **Safe-for-Flight Aircraft, Maintained Following Appropriate Procedures** and Released for flight by a **Qualified Person**, Flown Following Appropriate Procedures, overseen by a Qualified Person in Adequately Protected Airspace.

In aviation the concept of an airworthy aircraft is a key component to safe outcomes. An aircraft is airworthy if:¹⁷

- It is in a state that conforms with its approved design; and
- Is in a condition for safe operation.

Airworthiness be split into two parts, initial airworthiness and continuing airworthiness which contain the following elements:

- Initial Airworthiness
 - Design of Aircraft
 - Manufacture of Aircraft
 - Certification of Aircraft
- Continuing Airworthiness
 - Determining if the aircraft has deviated from its airworthiness state
 - Undertaking maintenance to restore the aircraft to an airworthy state

4.2.1 Initial Airworthiness Concepts

Initial airworthiness comprises all the activities that result in an article (i.e. an aircraft or an aircraft component) being produced and verified as airworthy. These are further delineated into three responsibility components for; the **design**, **manufacture**, and **certification** of aircraft or aircraft components.

- Responsibility for the approved design of an aircraft or aircraft component

This is the entity that holds and maintains the design data that demonstrates compliance with the regulatory requirements. If design flaws are discovered, this entity is responsible to update/re-design to fix these flaws. As an example, a Type Certificate

¹⁶ See Regulation 101.022 of CASR 1998

¹⁷ See Regulation 42.015(2) of CASR 1998

Holder under Part 21 of CASR is responsible for the approved design mentioned in the type certificate.

It should be noted that not all aircraft are required to have a type certificate in order to be issued a Certificate of Airworthiness. Examples include Light Sport Aircraft (LSA) and Experimental Aircraft.

- Responsibility for the manufacture of an aircraft or aircraft component

The manufacturer is responsible for ensuring that the act of manufacturing an actual instantiation of the design complies with regulations and results in an aircraft that meets the approved design. In most cases the type design holder and the manufacturer will be the same organisation. Some examples of these approvals are; production certificates (PC), Australian Parts Manufacture Approval (APMA) or Australian Technical Standard Order Authorisation (ATSO-A).

Again, there are some exceptions to the generic requirement to have a manufacturing approval for an aircraft. LSA and Experimental Aircraft do not require a manufacturing approval to be issued a CoA.

- Responsibility for the ensuring manufactured aircraft or aircraft components comply with the design

An independent responsibility exists to ensure that the process of manufacture has successfully created an aircraft that conforms to the approved design and to ensure that the design itself does not introduce any issues that result in the aircraft not being safe for flight.

This is traditionally achieved through the issuance of a certificate of airworthiness, which an aircraft is required to have unless a regulation exists that removes this requirement,¹⁸ such as in Regulation 101.046 of CASR for RPA other than Large RPA. This certificate provides a reliable indicator of the airworthiness of the aircraft.

4.2.2 UAS and Certificates of Airworthiness


The figure below shows a flow chart to determine if a UAS needs a certificate of airworthiness:

At the current point in time, Large RPA (an aircraft with a maximum take-off weight of greater than 150kg or an airship with envelope capacity of more than 100m³)¹⁹ must obtain a special certificate of airworthiness in the Restricted or Experimental Category, which are categories developed for conventionally piloted aircraft that are operationally restricted due to not meeting the normal airworthiness standards for aircraft of similar types.

¹⁸ ss. 20AA(3) of the Civil Aviation Act 1988

¹⁹ See Regulation 101.022 of the Civil Aviation Safety Regulations 1998

Figure 7 - UAS Certificate of Airworthiness Requirement diagram



The Restricted Category is a regime intended for special purpose operations²⁰ of conventionally piloted aircraft (i.e. surplus military aircraft or normal aircraft modified for special purposes). This results in requirements that are effectively conventionally piloted aircraft standards minus the parts that these aircraft cannot meet (i.e. weight and balance due to additional payloads). For UAS that are extremely close in design and intent to conventionally piloted aircraft, this is not an unreasonable requirement, however for UAS that are extremely different to conventional aircraft, these requirements end up being extremely restrictive (i.e. a 5kg powered lift, battery powered aircraft does not compare to a 3000kg fixed wing conventional aircraft).

An Experimental Certificate on the other hand is designed for very low risk operations (i.e. for testing completely unproven aircraft or for testing small modifications like WiFi in large transport aircraft) and is contingent on the operational scenario within which the certificate is issued.

Both of these CoA categories were never designed to cater for UAS, and as such are not fit for purpose. They are overly restrictive and do not allow for the unique design considerations of UAS.

For aircraft that are not large RPA, there is no requirement to have a certificate or airworthiness, however there is a caveat that an aircraft without a certificate of airworthiness cannot operate over a populous area without one issued.²¹ The

²⁰ The list of special purpose operations is in Subregulation 21.025(2) of the Civil Aviation Safety Regulations 1998

²¹ Note that this only applies if the UAS is operating at a height below which, if an aircraft component fails, the aircraft is able to clear the populous area, see Regulation 101.280 of the Civil Aviation Safety Regulations 1998

airworthiness of these platforms is catered for under a concept of “operational airworthiness”, where the operational approval considers the design and reliability of the system in context to the risk of the proposed operation. This is done on a case-by-case basis and each new operational approval must go through repeated analysis to demonstrate a safety case.

In general, it can be said that the removal of the requirement of a Certificate of Airworthiness for non-large RPAS is both a blessing and a curse for the UAS industry. Whilst it does increase operational flexibility due to the stringent nature of current, conventionally piloted airworthiness standards, it also increases the complexity of demonstrating airworthiness compared to having a certificate to demonstrate that an aircraft is airworthy.

4.2.2.A Continuing Airworthiness Concepts

For any aircraft that is operated, operational wear and tear, discrete source damage etc. will affect it and cause it to deviate from its approved design. There is a responsibility to ensure that the aircraft continues to be airworthy throughout operation, which will usually sit with the owner/operator of the aircraft. There are two main responsibilities in continuing airworthiness;

- **Ensuring the continuing airworthiness** of the aircraft by determining when/if the aircraft should undertake maintenance (scheduled maintenance or inspection-based maintenance); and
- **Certification of Maintenance**, which is required to demonstrate the processes used to bring an aircraft back into an airworthy state has been completed, and has been undertaken correctly.

Both of these aspects are required to ensure that an aircraft continues to be airworthy.

4.2.3 UAS and Continuing Airworthiness

For UAS, ensuring the continuing airworthiness of the aircraft is covered under the ReOC requirements in Regulation 101.340 of CASR 1998. This requires a maintenance controller role within the organisation that has responsibilities to ensure that all maintenance carried out on the operator’s RPA systems is carried out in accordance with the operator’s approved documented procedures.²²

For large RPA, they must be maintained as a Class B aircraft, which is a class contained within the older Civil Aviation Regulations 1988 (CAR). As large RPA can only have a CoA issued in the restricted or experimental categories, this restricts them to complying with Regulation 42CA and 42CB of CAR for their maintenance schedule.

4.2.4 UAS and certification of maintenance

For the certification of maintenance undertaken, this too is handled by the maintenance controller as stated in CASA AC 101-05. For cases of large RPA, any

²² CASA AC 101-05, as legislated by Subregulation 101.340(c) of the Civil Aviation Safety Regulations 1998 <https://www.casa.gov.au/sites/default/files/ac10105.pdf>

person certifying the maintenance must have a maintenance authority²³ issued to them in order to certify the maintenance undertaken.

Key Regulatory Opportunities for Highly Automated/Autonomous Systems

This section highlights some key opportunities in the airworthiness realm to increase the adoption and use of highly automated/highly autonomous systems.

AIRW.01 Providing an optional CoA route to demonstrate safety of UAS design and manufacture at different risk levels

Regardless of the level of autonomy of a system, a legal mechanism to recognise the design integrity, reliability and robustness of a UAS at varying risk levels with accompanying guidance on how to demonstrate compliance to a regulator will increase the ability of manufacturers to sell recognised products and reduce the time it takes for operators to obtain approvals.

An example of this method would be the US Federal Aviation Administration's new "Durability and Reliability" Type Certificate. This method allows lower risk UAS and operations to demonstrate reliability and robustness through operational and functional tests. A key concept in this method is a scalable suite of testing requirements dependent on the risk of the operation.

Instead of requiring a CoA for all operations, the provision of an **optional** choice to obtain a certificate of airworthiness in a UAS category at a given risk level (i.e. a SAIL III certificate of airworthiness under the JARUS SORA methodology) would allow the flexibility of not obtaining a CoA (i.e. undertaking a first principles analysis of safety of the aircraft and operation) if that does not align well with the operations objectives (i.e. R&D), and also the certainty of a recognition of an aircraft design through a CoA if desired.

AIRW.02 Updating CoA requirements for Large RPA

The requirement for a CoA for Large RPA could be updated to reflect a truly risk based regulatory regime. A Large RPA could be operated in an area within which the risk is extremely low (i.e. flight tests within a remote flight test range), and should not require a CoA. Conversely operating a Large RPA over a populous area is a high-risk operation and would require an in-depth analysis and certification to demonstrate safety.

By aligning the requirements for Large RPA to that of all RPA as depicted in AIRW.01 this would reflect a true risk based regulatory framework for the initial airworthiness of all UAS platforms.

AIRW.03 Adaptable CoAs

The conventionally piloted aircraft initial airworthiness process is a very static process. Once the design is approved, it is "frozen" in time, meaning any updates or changes must go through a significant process to be incorporated into an updated design approval. For aircraft designed to last for over 20 years, this makes a lot of sense, however for UAS, which have much lower design objectives and a faster design cycle, a method allowing reliability growth of a design over time is needed. Additionally, for adaptive learning systems, an adaptive certification regime is necessary to allow these types of systems to exist in the civil aviation realm.

²³ FAA, "Certification for Advanced Operations Unmanned Aircraft Systems (UAS) – Criteria for Special Classes."

4.3 A discussion on Airspace Integration regulatory requirements for complex UAS operations

Safe Operations [occur in] use Registered (when needed) and Safe-for-Flight Aircraft, Maintained Following Appropriate Procedures and Released for flight by a Qualified Person, Flown Following Appropriate Procedures, overseen by a Qualified Person in **Adequately Protected Airspace.**

Airspace management, as outlined by ICAO Doc 9854,²⁴ can be described using three layers of conflict management:

- Strategic conflict management through airspace organisation and management;
- Separation provision; and
- Collision Avoidance.

At some level, the rules associated with airspace classes and flight rules are all aimed at reducing the risk of a collision between two aircraft through these three layers.²⁵ This section discusses the following components of airspace management:

- Flight rules associated with aircraft operations (i.e. flying by visual flight rules or instrument flight rules) and how this affects rules and equipage requirements, including a discussion on non-controlled aerodromes.
- Airspace classes and their effect on the airspace organisation, rules, ATC involvement, and equipage.
- UAS operations in airspace and the issues related to not complying with the current rules associated with current airspace management.

4.3.1 Conventionally Piloted Aircraft – Visual Flight Rules and Instrument Flight Rules

For conventionally piloted aircraft, there are two types of environmental conditions under which flights can take place, Visual Meteorological Conditions (VMC) and Instrument Meteorological Conditions (IMC). Visual Flight Rules (i.e. day and night VFR) can only occur under VMC conditions (unless there is permission to undertake special VFR flights), whilst Instrument Flight Rules (IFR) can be undertaken in any meteorological conditions. Associated with these flight rules are minimum equipage (as detailed in Civil Aviation Order 20.18), operational procedures, licencing and training requirements.

²⁴ ICAO, “Global Air Traffic Management Operational Concept.”

²⁵ Note that although ICAO Doc 9854 is specifically referring to the Air Traffic Management concept, these three layers may still exist (i.e. radio carriage and use within a Common Traffic Advisory Frequency (CTAF) is a form of separation provision for uncontrolled airspace)

This is important because IFR aircraft are required to carry certain equipment (in particular Automatic Dependant Surveillance – Broadcast or ADS-B)²⁶ that make these aircraft electronically conspicuous. For VFR flights, whilst they can elect to carry ADS-B or other equipment that IFR aircraft carry, they are not required to and are generally not electronically conspicuous, being labelled as “non-cooperative” aircraft.

4.3.2 Airspace Classes

Under the Airspace Regulations 2007, CASA has classified airspace into one of seven Airspace Classes (A, B, C, D, E, F or G). In addition to the VFR and IFR rules, these airspace classifications delineate rules and responsibilities for aircraft and Air Traffic Control. A key differentiator within airspace classes is whether they are controlled or not. Airspace classes A-E are considered controlled as ATC will provide some form of service to IFR and VFR traffic (dependent on the class of airspace and practicability). Class G airspace is considered “uncontrolled”, however ATC may still provide services if practicable. More information on the different requirements in classes of airspace based on whether the operation is VFR or IFR can be found [here](#).

4.3.3 Aerodromes

Aerodromes are areas of high risk for mid-air collisions, due to the nature of traffic flow and the increased density of aircraft movements being funnelled at, through or from these aerodromes. Additional regulatory requirements have to be met in order to maintain an acceptable level of safety in these environments. Controlled aerodromes are those that occur in controlled airspace (i.e. Class C and D aerodromes), and have rules and equipage requirements associated with the airspace class. Non-controlled aerodromes are those that are in Class G airspace and may require the mandatory carriage of radio. Further information on non-controlled aerodromes can be found [here](#).

4.3.4 UAS and airspace integration

The entirety of these airspace classes, flight rules, operational procedures (ATC clearance, joining the circuit etc.), training, licencing, and minimum equipage requirements (radio, ADS-B etc.) have been developed around preventing mid-air collisions between aircraft that currently exist in the airspace. When you introduce UAS – which are not equipped, do not fly, and cannot see-and-avoid like other aircraft – you introduce additional risks into the airspace.

Under Regulation 101.020 of CASR, UAS are exempt from the requirements described above, effectively removing the requirement to fly VFR or IFR. Instead RPAS are restricted to operating within VLOS unless they are approved to operate BVLOS (as described in 0). In order to obtain an approval, it needs to be shown that all residual airspace risks that arise due to the operation of the UAS have been appropriately managed. Most UAS are not able to fully comply with the current airspace system rules and requirements, some examples of this include:

²⁶ 9B.8 of Civil Aviation Order 20.18

- Being unable to properly integrate into Class G airspace, as VFR aircraft rely upon the rules of the air, particularly “see and avoid”, for collision avoidance. UAS that cannot meet the intent of this requirement, increase the residual risk of collision in that airspace.
- Being unable to carry certificated equipment for IFR flight, meaning that ATC cannot apply standard separation procedures to these aircraft (i.e. an ADS-B system below the required Software Integrity Level is invisible to ATC and hence separation provision between aircraft in the relevant airspace types cannot be undertaken).
- For UAS that rely upon a robust command, control and communication (C3) link for radio broadcast, losing C3 to the UAS adds an additional failure path preventing the use of radio.
- That UAS may not have the flight performance for conventionally piloted aircraft, resulting in issues where separation standards cannot be appropriately applied or followed (for example HALE aircraft may fly so slow that their groundspeed is negative. There is no separation standard that applies for an aircraft that can travel backwards so ATC cannot perform their function correctly).

This results in excessive restrictions on BVLOS operations, typically requiring a demonstration that the operational volume is sanitised of other airspace users. This can be done using various techniques:

- Significant stakeholder engagement to determine who is using the airspace in the area of operations, when they fly and how to avoid interacting with them (by operating **when** they are not there, or **where** they do not operate).
- Restricting operations to areas that conventionally piloted aircraft cannot or would not go (i.e. shielded operations below the height of buildings, operations at night below the minimum safe altitude).
- Obtaining a restricted area for the operation, effectively sanitising the airspace of other airspace users without permission of the person responsible for the restricted area.

While these methods do facilitate some accommodated access to airspace, these cannot scale and are not cost or time effective for the operator or the regulator.

Key Regulatory Opportunities for Highly Automated/Autonomous Systems**AIRSPC.01. Being able to comply with the intent of remaining well clear and “see and avoid”**

A key part of airspace integration is being able to avoid collision hazards with other aircraft. Being able to remain well clear, follow the right of way rules and see-and-avoid other aircraft is predicated on the ability to see out of the cockpit.

For highly automated/autonomous aircraft flying BVLOS operations, this results in being unable to properly integrate into Class G airspace. High level objectives, regulations and standards to demonstrate that a UAS can meet the intent of these rules, i.e. through “sense-and-avoid” would allow significantly more access to complex airspace for UAS operations and reduce the burden on both operators and CASA to undertake resource intensive reviews of strategic mitigations to reduce airspace encounters to an acceptable level.

AIRSPC.02. UAS specific airspace classes (Class “UTM”, drone delivery corridors etc.)

As a key initiative in the National Emerging Aviation Technologies policy statement,²⁷ iterating towards a UAS Traffic Management concept will assist in the integration of UAS into low level airspace. There is an opportunity to create a “UTM” class of airspace that would have specific equipage requirements for any aircraft operating within them, such that UAS would easily be able to separate from each other and conventionally piloted aircraft that operate within this airspace.

There are also opportunities for other ways to increase access to airspace without a need for significant airspace class changes. Adopting drone corridors (i.e. restricted or danger areas) that are labelled and intended for drone use in rural areas (i.e. between towns) would allow UAS to operate BVLOS within reasonable bounds without extensive effort for stakeholder engagement and strategic mitigations to reduce the risk of collision.

²⁷ Department of Infrastructure, Transport, Regional Development and Communications, “National Emerging Aviation Technologies Policy Statement.”

5 Aircraft operating for Defence (State Aircraft)

As exempted from the Civil Aviation Act, Australian defence aircraft are not regulated by CASA, unless specifically required within the regulations.²⁸ As these operations are undertaken by the department of defence in some capacity, there exists,

“moral and legal obligations to ensure risks to the health and safety of personnel arising from military aircraft operations are eliminated or otherwise minimised so far as is reasonably practicable. The basis for these safety obligations is found in Commonwealth Work Health and Safety Legislation and Common Law.”²⁹

Which are managed by the established Defence Aviation Safety Authority (DASA), the Defence Aviation Safety Regulations (DASR), and via the Defence Aviation Safety Program (DASP). We will be focusing on the DASR as a lens to what is required when operating under the DASR.

5.1 The DASR, Applicability, Structure and Decomposition of Aviation Activities

The DASR is enabled by the Basic Regulation, which sets the applicability of the regulations in Regulation 15 of the Basic Regulation (BR.015):³⁰

- (a) Aircraft, including any installed product, part and appliance, which are:
 - 1. owned by Defence;
 - 2. operated exclusively for or on behalf of Defence;
 - 3. the subject of statutory airworthiness responsibilities placed on Defence by the national civil airworthiness authority;
 - 4. designed or manufactured by an organisation for which the Authority ensures safety oversight;
 - 5. foreign military aircraft operating within national airspace shall comply with this Regulation.
- (b) Personnel involved in the operations of aircraft referred to in BR.15.A shall comply with this Regulation.
- (c) Operations of aircraft referred to in BR.15.A shall comply with this Regulation.
- (d) Defence aerodromes, including equipment, shall comply with this Regulation. Personnel and organisations involved in the operation of these aerodromes shall comply with this Regulation.
- (e) By way of derogation from BR.15.D, the Authority may decide to exempt from the provisions of this Regulation a Defence aerodrome with only a few movements per year.
- (f) Defence ATM/ANS shall comply with this Regulation. Systems and constituents, personnel and organisations involved in the provision of ATM/ANS shall comply with this Regulation.

²⁸ See Section 4 of the Civil Aviation Act 1988 (Cth.)

²⁹ Department of Defence, “DASR - Aviation Safety Management System.”


³⁰ Department of Defence, “Defence Aviation Safety Regulations (DASR).”

The DASR regulation has been set up to regulate the following types of operations:

- Non-Defence Registered Aircraft
- Special Purpose Operations
- Flight Tests
- Air Cargo Delivery
- UAS Operations

Of which we are interested in UAS operations. DASR has been developed to align with the European Military Aviation Requirements (EMARs), which are military versions of EASA regulations. This results in an alignment to the Open, Specific and Certified UAS operation types used by EASA. The operational categorisation of the DASR is shown below:

Figure 8 - Defence Aviation Activity Delineations



5.1.1 The Open Category

The Open Category considers “low risk” UAS operations and contains all UAS operations that are effectively segregated from other airspace users and ground users and is restricted by weight class (the DASR includes three categories, Micro [$<100\text{g}$], Very Small [$0.1 \leq \text{MTOW} < 2\text{kg}$], and Small [$2 \leq \text{MTOW} < 25\text{kg}$]).³¹ As the mass of the aircraft increases across the Micro, Very Small, and Small categories, the operational restrictions increase commensurate to the increased risk to other airspace and ground users.

5.1.2 The Specific Category

The specific category is intended to cater for operations that are “medium risk” and are intended to be approved via single operating approval only (unlike the certified category, which requires multiple certifications for approval), based upon the results of a specific risk assessment for the operation. Under the DASR, specific category operations either:

- Align with a Standard Scenario defined within UAS.35 of DASR
 - Standard Scenario for Micro UAS (DASR 35(a))
 - Standard Scenario for Very Small UAS (DASR 35(b))
 - Standard Scenario for Defence Areas and Exercise Areas (DASR 35(c))
 - Standard Scenario for High Seas (DASR 35(d))
 - Standard Scenario for Trials and Experimentation (DASR 35(e))
- Are approved under a UAS operating permit (UASOP)

Standard scenarios are pre-defined operations where the regulator (DASA in this case) has undertaken a risk assessment of the assumed operation and UAS category combination, and provided operational, organisational, and technical requirements to be met to mitigate the risk SFARP.

For operations that do not fall under a Standard Scenario, a UASOP must be obtained, and a specific risk assessment for the operation must be undertaken and all controls and mitigations must be applied SFARP. This pathway requires significantly more effort than following a standard scenario, but does provide infinitely more flexibility. Guidance for risk controls for UAS that can be applied through a specific category risk assessment can be found in DASA Advisory Circular AC 001/2018.³²

5.1.3 The Certified Category

The certified category exists for “high risk” operations. These are generally considered operations where the risk is at least equivalent to conventionally piloted aviation. DASR specifically mentions that the certified category contains operations that wish to operate in all airspace classes and over all populous areas (similar to the expected operational freedoms of a conventionally piloted aircraft) Certified operations must:

³¹ See UAS.40 of DASR for Open Category classifications and associated operational restrictions

³² DASA, “Risk Controls for UAS Operations.”

- Be type certified in accordance with the initial airworthiness provisions of DASR Part 21
- Comply with all continuing airworthiness requirements in DASR (i.e. DASR Part M for continuing airworthiness management, DASR Part 145 for maintenance organisations and DASR Part 66 for maintenance personnel licencing)
- Operated under a Military Air Operator Certificate (MAOC)
- Comply with the standard rules of the air (DASR Part AO and DASR Part RoA)
- Be operated by a qualified military pilot (or determined equivalent set of requirements)

If we compare this to the single UASOP that must be obtained in the specific category, the certified category requires multiple approvals across the stakeholders involved in an RPAS operation at higher levels of rigour.

5.1.4 Differences between the DASR and CASR

Whilst DASA and CASA strive to align as much as possible, there are underlying requirements and objectives that cause some deviation between the two regulators and regulations. Key amongst the reasons are:

- The use of SFARP as the guiding principle for acceptance of risk; and
- The ability to take into account urgent safety matters, emergencies and compelling operational imperatives when applying the SFARP principle.

These two differences will be discussed in the next sections.

5.1.5 Work Health and Safety and the So Far As Is Reasonably Practicable (SFARP) Principle

The driving force behind the DASR is WHS legislation. The Australian Work health and safety legislation utilises the concept of ensuring that risks are eliminated or reduced, “so far as is **reasonably practicable**”,³³ which differs from the safety basis of the Civil Aviation Safety Regulations, which require that an operation does not create an adverse effect on the safety of air navigation (paraphrasing the requirements for issuing approvals under Part 11 of CASR 1998).

Reasonably practicable is defined in Section 18 of the Work Health and Safety Act 2011, reproduced below:

18 What is *reasonably practicable* in ensuring health and safety

In this Act, ***reasonably practicable***, in relation to a duty to ensure health and safety, means that which is, or was at a particular time, reasonably able to be done in relation to ensuring health and safety, taking into account and weighing up all relevant matters including:

- (a) the likelihood of the hazard or the risk concerned occurring; and

³³ Commonwealth of Australia, “Work Health and Safety Act 2011.”

- (b) the degree of harm that might result from the hazard or the risk; and
- (c) what the person concerned knows, or ought reasonably to know, about:
 - (i) the hazard or the risk; and
 - (ii) ways of eliminating or minimising the risk; and
- (d) the availability and suitability of ways to eliminate or minimise the risk; and
- (e) after assessing the extent of the risk and the available ways of eliminating or minimising the risk, the cost associated with available ways of eliminating or minimising the risk, including whether the cost is **grossly disproportionate** to the risk.

Key to the concept is the assessment of cost compared to risk, with the measure being whether the cost of elimination or minimisation is grossly disproportionate to the actual risk at play.

Usually, the SFARP principle is depicted as a triangle of decreasing risk. There are two regions, the intolerable region (whereby no amount of risk mitigation or elimination can lead to acceptance of the risk, and the SFARP region, whereby a risk can be accepted if the SFARP principle has been applied correctly.

We can contrast this process with the acceptable level of safety (ALOS) or target level of safety methodology (which is approximately what is used by CASA to measure “not adverse effect on the safety of air navigation”). In the ALOS methodology, once you reach a certain level of risk mitigation, you can “stop” implementing mitigations, because the risk has been minimised to an acceptable level. Now this does not mean that ALOS is inherently worse or better than the SFARP methodology, and they both have benefits and weaknesses. For example, although the civil aviation sector uses an ALOS like process, the acceptability of risk is so low that it may inherently meet or exceed the risk reduction in an SFARP process for another industry. This could be thought of as the intolerable region for the aviation ALOS being extremely low in the SFARP region shown below.

Figure 9 - The SFARP





Figure 10 - Comparing the TLOS and SFARP Processes




5.1.6 DASA's ability to accept higher levels of risk in the case of urgent safety matters, compelling operational imperatives or emergencies

In civil aviation, safety is the **most** important consideration³⁴, which results in the ALOS process in CASA being fairly rigid and immovable. This makes sense in the civilian case, however, under a defence context, flexibility is needed to allow for the various needs of defence that exist alongside safety. Within the DASR regulations this flexibility is created through Subregulation 3(g) of DASR BR, which states that the DASA regulations “Enough flexibility” to address special circumstances, such as urgent safety measures, and compelling operational imperatives or emergencies.

Whilst this is allowed, the scope of utilising this flexibility is heavily oversighted to ensure appropriate use. This gives DASA the flexibility to operate to both safety and other objectives as needed, whilst still applying the SFARP principle. We can view this on our previous SFARP vs TLOS graphic as the ability for defence to change the intolerable region depending on the operational context, in the figure below this is shown as an increase in acceptable risk (shown as $\Delta Risk$) due to a compelling operational imperative.

It should be noted that during normal circumstances, DASA does attempt to align as much as possible with the civilian safety standards and requirements that CASA and ICAO³⁵.

Figure 11 - How DASA can change its acceptance of risk depending on the operational imperative (example only)



³⁴ ss. 9A(1) of the Civil Aviation Act 1988

³⁵ See Subregulation 3(c) of DASR BR and Regulation RoA.05 and GM RoA.

6 Comparative studies – other countries integration of UAS

Integration of autonomous systems is a worldwide challenge. Every nation is on a unique path towards integrating highly automated and autonomous systems into their national airspace systems. Some nations are ahead in terms of access to airspace, some ahead in certification of systems, and others ahead in taking the front foot on adopting and accepting the use of data-driven systems.

This section provides comparative studies of international aviation authorities approaches to regulating UAS operations. We will focus on the two largest players in the regulation of aviation activities, the European Union Aviation Safety Agency (EASA) and the Federal Aviation Administration (FAA) of the United States.

6.1 European Union – European Union Aviation Safety Agency

6.1.1 How UAS are regulated by EASA

EASA is an agency under European Union and as such, covers the safe regulation of aviation activities that have the ability to affect over 400 million people. EASA's high level policy for the regulation and oversight of UAS operations has been driven by categorising UAS operations into three distinct categories:

- **Open Category**³⁶ – those operations considered low risk, with no operational approval is required before flight. Divided into three subcategories (A1, A2 and A3), based on weight and operational risk, it is somewhat akin to the Australian “excluded” category, however the weight and operational limits are broader within this open category compared to the excluded category and do not require notification to the national authority in order to commence.
- **Specific Category**³⁷ – those operations considered medium risk. If the operation is covered by a standard scenario, a CONOPS with pre-defined requirements and standards to be met, the applicant only needs to declare to the national authority that they will undertake the operation in accordance with the standard scenario (similar to the notification requirements for excluded category operations in Australia).

If the operation is not covered by a standard scenario, an operational authorisation must be obtained from the national authority through the use of an operational risk assessment that the operation is safe, with the SORA being an accepted means of compliance for undertaking the risk assessment. The specific category is very similar to “included” RPA operations (including autonomous aircraft) under Australian legislation.

At higher levels of risk (SAIL V and SAIL VI within SORA terminology), EASA requires that a type certificate is obtained for the system in order to meet the technical requirements under the SORA scheme.

At medium levels of risk (SAIL III and SAIL IV), EASA requires a process of “design verification” for the aircraft system on top of the operational authorisation.

³⁶ EASA, “Open Category - Civil Drones | EASA.”

³⁷ EASA, “Specific Category - Civil Drones.”

- **Certified Category**³⁸ – those operations considered intrinsically high risk. These operations will require certified air operators and certificated aircraft in order to operate (effectively the same requirements as conventionally piloted aircraft). This category is still under development, however EASA has stated the following three types of operation will be considered within the certified category:
 - International flight of cargo UAS under IFR
 - Drone operations in urban or rural environments using U-space services
 - Urban air mobility operations (i.e. air taxis with persons onboard)

Australia does not have an equivalent category to the European certified category at this point in time (although the Australian rules requiring a CoA for Large RPA could be considered a certified category, it does not necessarily relate to risk of the operation).

6.1.2 Progressive activities to further incorporate UAS and autonomy into aviation

As a large aviation authority with significant international regard, EASA is viewed as a leader of world class rules and many countries replicate the regulatory style of EASA's UAS rules. It is important to note that within the EU each country still has their own national aviation authority (i.e. the French DGCA, in the Netherlands, the NLR etc.). As a result, there exists competencies that EASA handles as an EU agency and competencies that individual countries NAA's also handle, leading to a complex regulatory picture.

EASA has taken the front foot when dealing with artificial intelligence and autonomy. In February 2020, they released the EASA AI Roadmap,³⁹ a document aimed at establishing a vision for integration of AI into aviation. Crucially this roadmap starts from an ethical baseline aligning with the EU Ethics Guidelines on Artificial Intelligence.

In April 2021, EASA released for consultation useable guidance for level 1 machine learning applications⁴⁰ (i.e. systems that provide assistance only to humans), introducing the four key pillars for incorporating ML into aviation. The pillars are:

- Trustworthiness Analysis;
- Learning Assurance (a commensurate process to software development assurance but tailored to the use of data driven methods rather than model based (deterministic, algorithmically specified);
- Explainability Analysis; and
- Safety Risk Mitigation.

This EASA document is intended to be a “Swiss Army knife” document, providing a starting point from which to begin understanding the analysis of these ML systems.

³⁸ EASA, “Certified Category - Civil Drones | EASA.”

³⁹ EASA, “EASA Artificial Intelligence Roadmap 1.0 Published.”

⁴⁰ EASA, “EASA Releases for Consultation Its First Usable Guidance for Level 1 Machine Learning Applications”; EASA, “EASA Concept Paper: First Usable Guidance for Level 1 Machine Learning Applications.”

6.2 United States – Federal Aviation Administration

6.2.1 How UAS are regulated by the FAA

The Federal Aviation Administration regulates commercial UAS under the following regimes:

- For small UAS within low-risk operations, Under Part 107
- For those that do not meet the requirements for Part 107, under conventionally piloted aircraft operational Parts (i.e. Part 91, Part 135)

Part 107 was introduced by the FAA in 2016 as a way to allow small UAS to operate without being bound by overly restrictive and onerous requirements for conventionally piloted aircraft, which were designed around the use of large, crewed platforms. Part 107 is restricted to aircraft that have a maximum take off weight of less than 55 pounds (approximately 25kg).

Part 107 starts off with an assumption that the UAS are not able to meet the definition of airworthy⁴¹ (i.e. there is no approved design to conform too), and hence operational restrictions are applied to prevent these operations from flying over persons⁴² and restrict operations to within visual line of sight, such that they can see and avoid other aircraft by yielding right of way. The operational limitations are somewhat similar to the requirements under the excluded category, however there are some discrepancies related to overflight, weight classes etc.

Under Part 107 the FAA may issue a waiver to the operational restrictions contained with Part 107. This process usually takes the form of a risk assessment to determine the additional compliance requirements needed in addition to Part 107 in order to maintain the safety air navigation.

In April 2021, the FAA updated Part 107 for operations over people (OOP), based upon recommendations from the Micro Aviation Rulemaking Committee⁴³ in 2016. This allows operations over people under four categories of operation:

- **Category 1** – UAS that weigh less than 0.55 lbs (250g) and do not contain any exposed rotating parts. These restrictions are based upon work done to identify the “harmless” threshold for UAS based upon blunt force trauma to humans upon impact.⁴⁴
- **Category 2** – UAS that cannot impart more than 11 foot-pounds (14.9J) of kinetic energy to a human during an impact and do not contain any exposed rotating parts. Demonstration of this capability is intended to be self-declared through compliance with industry consensus standards.
- **Category 3** – UAS that cannot impart more than 25 foot-pounds (33.9J) of kinetic energy to a human during an impact and do not contain any exposed rotating parts. Operations must not be sustained over people not directly

⁴¹ Department of Transportation, Federal Aviation - 14 CFR Parts 21, 43, 61, et al. Operation and Certification of Small Unmanned Aircraft Systems.

⁴² Unless they meet the requirements stipulated in regulation 107.39 of CFR 14

⁴³ FAA, “Certification for Advanced Operations Unmanned Aircraft Systems (UAS) – Criteria for Special Classes.”

⁴⁴ la Cour-Harbo, “Mass Threshold for ‘Harmless’ Drones -.”

involved in the operation (although it may be transit over people not involved in the operation) and can only occur within or over a restricted access site. Demonstration of this capability is intended to be self-declared through compliance with industry consensus standards. In both Category 2 and 3, the energy thresholds selected are based on conservative values for causing injuries to humans based upon blunt force criteria.

- **Category 4** – small UAS that do not meet Categories 1, 2 or 3 are required to have a certificate of airworthiness with appropriate operational limitations and procedures to ensure safety of persons overflown.

In order to operate over open-air assemblies of people (for categories 1, 2 and 4), operations must comply with Remote identification standards covered in Part 89 of FAA regulations.

For operations that do not meet the requirements of Part 107, they are to be operated under conventionally piloted aircraft operational parts. In the current state, this results in significant operational restrictions with accompanying exemptions (i.e. operating without a certificate of airworthiness, operating BVLOS without being able to comply with see and avoid requirements).

6.2.2 Progressive activities to further incorporate UAS and autonomy into aviation

The FAA is critically aware of the integration issues for UAS operations that cannot meet the requirements of Part 107 (including operations that can occur under waiver). To that end, the FAA undertook Integration Pilot Programs (IPP) between 2017-2020, which aimed to demonstrate technology and operations that will lead to future integration into the National Airspace System. Whilst the IPP has come to an end, a secondary program, “UAS BEYOND” continues these partnerships and focuses on three key challenge areas. These are:

- Scalable and repeatable BVLOS operations;
- Quantifying societal benefits of UAS use; and
- Identifying and understanding societal concerns with the use of UAS.

To further reduce the regulatory burden of UAS operations under conventionally piloted rules, the FAA have introduced a new method for obtaining a type certificate for lower risk UAS, called “durability and reliability”.⁴⁵ For traditional crewed aircraft, the certification safety requirements for aircraft are such that there is not a possibility to utilise operational reliability as claims against these requirements (i.e. in order for a large transport category aircraft to demonstrate a 1 in 1 billion hour catastrophic failure rate with operational data would require 3 billion hours of operation, or over 340,000 years of operation), and instead design, analysis simulation, and functional testing are used to show that these requirements are met.

For lower risk operations, that have lower reliability requirements, it is completely feasible to use operational demonstration testing to show that safety targets are met (i.e. to demonstrate a hazard rate of less than 1 failure per 10,000 hours would only

⁴⁵ FAA, “Certification for Advanced Operations Unmanned Aircraft Systems (UAS) – Criteria for Special Classes.”

need 30,000 hours, or just under 3 and a half years).⁴⁶ This is the drive for the new “durability and reliability” type certificate process for UAS operations. There are at least ten different UAS manufacturers using this method to obtain type certificates and subsequent certificates of airworthiness in order to operate under conventionally piloted aircraft rules.

A final key element of the FAA’s integration of UAS into airspace is research and development, including:

- The Alliance for System Safety of UAS the Research Excellence (ASSURE)⁴⁷, a conglomerate of research institutions attempting to understand various facets related to safety in the context of UAS, including both System Safety and Autonomy for UAS; and
- A collaboration with NASA on an UAS Traffic Management concept,⁴⁸ which has driven a substantial amount of current work in proto-UTM systems (i.e. the Flight Information Management System request for proposal by Airservices Australia⁴⁹); and
- UAS Test Ranges, of which the US has seven operational at this point in time, allowing for both non-commercial research and commercial CONOPS development/demonstration.

⁴⁶ This is for a single aircraft, for 100 UAS operating constantly to demonstrate reliability, this target of 30,000 hours would only take about 125 days of operation (assuming no failures during testing and a binomial distribution)

⁴⁷ Assure UAS, “Assure - THE FAA’s Center of Excellence for UAS Research.”

⁴⁸ NASA, “UAS Traffic Management (UTM) Project.”

⁴⁹ Airservices Australia, “Flight Information Management System (FIMS): Request for Proposal | Engage Airservices.”

7 Current methods for assurance of highly automated/autonomous UAS (JARUS SORA)

The previous subsections have been discussing the regulatory frameworks and accreditation requirements that exist for UAS operations in Australia. From that foundation, the next crucial understanding is how to demonstrate compliance with those higher-level regulatory objectives, and this brings us to assurance. We are interested in the concept of assurance with respect to aviation safety claims. CASA's expectation for the issuance of approvals for complex, highly automated and autonomous operations is that there is no adverse effect on aviation safety.

When attempting to demonstrate a claim like, "this UAS operation does not have an adverse effect on aviation safety", we are trying to ascertain whether the evidence provided for this claim is "enough" to reduce the uncertainty around the truth of this statement. There exists some threshold that defines "enough" certainty (or conversely, the maximum allowable uncertainty) to be considered acceptable to meet this claim using evidence. We could say that the safety claim can be considered true if the probability that the safety claim is not true given the evidence must be less than a maximum allowable uncertainty

$$1 - P(\text{Safety Claim} | \text{Evidence}) \leq \text{maximum allowable uncertainty}$$

We might be able to state that assurance is the acceptable process by which enough appropriate evidence is provided to meet this minimum standard of uncertainty. Now in reality it may not be clear or possible to attach numbers to any of these variables.

Within the aviation safety context, it is important to know the following concepts to drive what the appropriate evidence is to demonstrate our claim "this operation does not have an adverse effect on aviation safety":

- how risk management is understood and used in reducing the likelihood or consequence of hazards, particularly how engineering approaches tend to drive likelihood requirements and uncertainty using the assessed consequence of a hazard. Sections 7.1, and 7.2 detail the safety risk management process and system safety processes.
- How the measurement of risk can be done and the different considerations that apply for unpersonned aircraft. Sections 7.3 detail methods of measuring risk for aviation applications and how the tolerability of risk changes based on whether a person at risk is a first, second or third party to the operation.
- the acceptable level of safety for current aviation operations and how this is applied for different operations using the concept of the aviation safety continuum. This is detailed in Section 7.4.
- a model for how uncrewed aircraft operations can casually lead to consequences for third parties (in the air and on the ground). Sections 7.8.1 and 7.8.2 detail risk models for ground and air risk respectively, and Section 7.9 details how the bow-tie process can be used to understand the different causal events that lead to top-level events and realised hazards.
- an acceptable depth of analysis process based upon the risk a UAS operation poses to third parties in the air and on the ground. Section 7.11 introduces the JARUS Specific Operations Risk Assessment (SORA) methodology, which ties in all the previous assurance discussion with a risk proportionate, regulator acceptable process for demonstrating a reasonable safety case (including a depth of analysis process) based upon the risk an operation poses to third parties.

Although the SORA process is accepted internationally, this does not mean that all the issues and problems for highly automated and autonomous aircraft are solved and it is all clear skies ahead. Section 8 closes this loop by highlighting the existing gaps in current best practice assurance processes used in the aviation industry, particularly those issues that affect the true integration for autonomous aircraft and identifies potential solutions to these issues.


7.1 Safety Risk Management

We are all intimately familiar with the concept of risk as a combination of likelihood and consequence:

$$\text{Risk} = \text{likelihood} \times \text{consequence}$$


This results in a graph shown in the Figure below, with risk isopleths (lines of constant risk) highlighted:

Figure 12 - Classic Risk Curve



Normally the risk graph is viewed on a log-log curve:


Figure 13 - log-log curve of risk



Now in theory this continuum of risk could be useful for defining limits to the combinations of likelihood and consequence, however consequences are not easily definable along a continuum. Instead, this graph is bucketized into regions of discrete

likelihood and consequences, resulting in the familiar tables found in most risk assessment processes:

Figure 14 - How the risk continuum is replicated in risk tables



In risk management practices, the ability to analyse and categorise risk into buckets (such as “extreme”, “high”, “moderate” etc.) allows the user to then apply a known set of risk management tools based upon this level. Within each bucket of increasing risk the allowable uncertainty of any safety claim is driven by the risk of the operation or system. The higher the risk, the lower acceptable uncertainty for any safety claims, and the higher the evidentiary requirements needed, for example:

- more detailed analysis of a hazard and hazard controls,
- wider coverage of potential failure conditions (i.e. including outside expected operating conditions, common cause analyses)
- a quantitative analysis complementing qualitative analyses
- specific architecture requirements (redundancy, independency, resilient design etc.)

For example at “low” levels of risk, it’s likely that nothing more needs to be done to manage the risk. Meanwhile, at “extreme” levels an exorbitant amount of work must be done to demonstrate that the risk has been managed appropriately. Effectively as the risk increases, the amount of uncertainty allowed in the determination that the risk has been managed decreases.

Alongside this higher level of evidentiary requirements, most risk management processes require higher managerial oversight to accept a higher given risk. This ensures that the right management level has knowledge and responsibility for ensuring the risk has been acceptably managed.

Safety risk management is a huge topic in aviation, and particular resources that may be of use include:

- [Federal Aviation Administration “Safety Risk Management” webpage](#), which has significant information and tools that can be applied to the aviation context
- [CASA’s Safety Management Systems webpage](#), which contains guidance on how to implement a safety management system, which is a required system for some aviation operations.


7.2 Systems Safety Engineering in Aviation

In aviation system safety engineering, a process that looks at aircraft systems and their effects on safety, this process is modified slightly. Using the risk equation, the maximum allowable probability of an event type (i.e. catastrophic, hazardous, major) is set based upon the allowable risk isopleth and the consequence of that event:

$$Likelihood \leq \frac{Allowable Risk}{Consequence}$$

Where the allowable risk is based upon some analysis of the societal acceptance of risk in different types of operations (see 0 for the aviation safety continuum) and the consequence is based upon the credible outcome of various hazards that aviation operations experience. An example of this is shown below:

Figure 15 - Aviation System Safety Acceptability Region



This allows for detailed safety assessment processes to be allocated based upon the severity of a realised hazard (Minor, Major, Hazardous, or Catastrophic), which is derived from the allowable failure condition likelihood (Probable, Remote, Extremely Remote, Extremely Improbable).

The system safety process assesses the consequence of a hazard against pre-defined word pictures found in guidance such as AC 23.1309-1E and then allocates qualitative likelihood requirements. These requirements must be met using allocated depth of

analysis and development assurance rigour, essentially providing the mechanism to ensure minimisation of uncertainty to an acceptable level based upon the hazard.⁵⁰

It is important to note that the depth of analyses requirements based on failure condition severity or failure condition likelihood depends on the type of operation being undertaken and the risk it poses. For instance, under Part 23 (Normal Category Aeroplanes), risk (and the subsequent requirements) is allocated by the number of seats onboard the aircraft. These severity definitions and requirements have been developed through decades of crashes, deaths, and loss, leading to an incredibly detailed set of requirements for conventionally piloted aircraft. For unpersonned aircraft, there is no risk to persons onboard and instead we need to understand the risk to second and third parties instead. This can result in risk levels that are substantially lower than that what occurs in conventionally piloted aviation, meaning that the in-depth requirements from the current aviation system safety process are not proportionate to the risk of these UAS operations. We need a way to modify the requirements to proportionate levels for UAS operations.

7.3 Measuring risk for unpersonned operations

In order to modify the requirements for risk, we need to begin by considering the definitions of risk to individuals to build a more generic picture of risk, then push this towards an acceptable level of safety for unpersonned aircraft.

7.3.1.A Individual Risk (risk to a specified individual)

Individual risk is usually described as a probability of a fatality occurring to an individual during a given event (i.e. probability that an individual dies per year, probability of an individual dying to per rocket launch, probability of an individual dying per aircraft per hour of exposure). We can write this individual risk due to event E as:

$$IR_i(E) = P(E) \times P(\text{indv } i \text{ affected}|E) \times P(\text{casualty}|\text{indv } i \text{ affected}, E)$$

Where $P(E)$ is the probability that an event E occurs over a given timeframe, $P(\text{indv } i \text{ affected}|E)$ is the probability that the i^{th} individual is affected by event E, on the condition that E occurs and $P(\text{casualty}|\text{indv } i \text{ affected}, E)$ is the probability that the individual experiences a casualty, given they have been affected by event E.

For example, the risk that an individual person dies in a year (utilising the 2018 Age-Standardised Rate from the Australian Bureau of Statistics⁵¹) is approximately 0.00327 (approximately 1 in 300. For a person aged 35-39, this risk is closer to 1 in 1200).

7.3.2 Societal/Collective Risk (risk to a population of people)

If we were to sum the risk to an individual over all the individuals at risk, we arrive at the collective risk to a population of persons, measured in fatalities per given event

⁵⁰ Note that recently Part 23 has been updated, and the Classes referred to in AC 23.1309-1E do not fully correspond to the new “level of certitude” that is based upon the number of seats onboard the aircraft

⁵¹ AIHW, “Deaths in Australia.”

(i.e. number of fatalities per year, number of fatalities per rocket launch, number of fatalities per aircraft per hour of exposure), usually referred to as the “expected casualty” or “expected fatality” rate (depending on whether you are looking at fatalities or casualties). We can represent the expected fatality rate (EF) as:

$$E_F(E) = \sum_{\forall i} IR(E)_i$$

This adds up the risk to each individual from event E. Using the ABS data on deaths per year, we can see that the expected fatality rate for 2018 was approximately 82,320 people/year. This is calculated by multiplying the individual risk by the number of people exposed to the risk, i.e. the entire population of Australia in 2018.

7.3.3 F-N curves (Frequency of events that cause more than N fatalities per event)

The previous discussion does not take into account societies’ risk aversion to catastrophic events. As an event increases in its consequence, the willingness of a person or persons to accept the risk decreases non-linearly to the point where there may be no willingness to accept an event occurring due to the consequence (i.e. the destruction of the world, death of all persons on Earth). F-N curves track the acceptability of the frequency of events that cause more than N fatalities per event. They effectively plot collective risk for events that cause differing fatalities per event.

In the next figure, as the number of fatalities per event increase, the frequency of the event drops commensurately until a point is reached where the risk of the event occurring cannot be justified at any frequency⁵².


Some examples of risk aversion in safety requirement include Section 4.3 of the Common Risk Criteria Standards for National Test Ranges – Supplement⁵³ or the lower allowable mid-air collision rates in controlled airspace (where passenger aircraft mid-air collisions can kill thousands of people).

While not usually considered by an applicant, risk aversion and societal acceptance of risk is something that regulators must contend with every day.

⁵² Porter and Morgenstern, “Landslide Risk Evaluation in Canada.”

⁵³ Range Commanders Council, “Common Risk Criteria Standards for National Test Ranges, Standard 321-20.”

Figure 16 - F-N Curve



7.3.4 Risk tolerability – Acceptance of risk for involved and uninvolved persons

Another confounding factor within this risk discussion is informed consent and acceptance of higher risk. There is an understanding and acceptance by both society and these individuals that a higher rate of fatality is acceptable under certain conditions. Examples of this in aviation include skydiving and flying in warbirds, which are commercial operations with known risks that involved parties need to acknowledge and accept before flight (see Division 132.C.2 of CASR).


Contrast this with persons who are uninvolved with an activity and have not had a chance to accept the risk of the activity. Society will generally try to protect these people to an extent that matches an expectation of safety for persons going about their everyday lives.

For the purposes of this analysis, we will only be considering third parties, i.e. those that have not directly been involved in or been part of the risk acceptance of UAS operations.

7.4 The Aviation Safety Continuum

CASA is committed to providing risk-based, proportionate regulation for aviation. This already exists in some form (although not perfectly) in conventionally piloted aviation and can be seen through the contrast in requirements for parachuting operations (Part 105 of CASR) and regular public transport aircraft (RPT). This can be visually represented as shown below:

Figure 17 - The aviation safety continuum (Note: position of CASR Parts on the continuum are for illustrative purposes only, actual positions along this theoretical curve will vary)



There exists a safety continuum for traditional manned aviation, that reflects the public's risk appetite for various activities. Notably, these traditional aviation types can only exist on the right-hand side of the dotted line in Figure 17. This Figure represents the point along the line at which, if the worst outcome of an operation occurs, at least one person will die (the pilot/parachutist in a single-pilot/solo-parachuting accident). Highly automated/autonomous aircraft can exist anywhere along the horizontal axis, from posing almost no risk at all (i.e. a small UAS flying in the middle of Antarctica), to posing as much (if not more) than a large transport category (CASR Part 121) aircraft.

7.5 Acceptable levels of safety for aviation

In order to answer the question, “what is the acceptable level of safety for aviation?”, one method can be accepting that the current state of aviation is already at an acceptable level. This would then imply that the current standards used in aviation should be appropriate for unpersonned aircraft operations. However, for truly unpersonned aircraft (i.e. no persons on board, regardless if crew or passengers), the only persons to consider at risk are those 2nd (i.e. involved in the operation but not onboard) and 3rd parties (uninvolved public) in the air and on the ground (as well as critical infrastructure). Current manned aviation standards are primarily focused on protecting persons on-board the aircraft. A secondary effect of this protection for those onboard is a high level of safety for third parties (i.e. safety of other aircraft and third parties on the ground).

We can crudely delineate accepted levels of safety for unpersonned aircraft into the following categories:

- Second party risks (those involved in the operation)
- Third party ground risk
- Third party air risk
 - in controlled airspace
 - in uncontrolled airspace

As we are specifically dealing with unpersonned vehicles, the only true second parties are those persons involved with the operation (i.e. ground/flight crew).

7.6 Acceptable levels of safety for third parties on the ground

Under the paradigm of “not less safe than manned aviation”, many sectors (aviation, nuclear, rocketry etc.) have derived acceptable risk to third parties on the ground due to aviation. Under the JARUS SORA model, the maximum allowable risk to third parties on the ground is less than 1 fatality per 1 million hours of exposure to aviation risk⁵⁴ (1×10^{-6} ph). Notably, for most populations at risk, meeting this requirement will automatically meet many individual risk targets that have been set in other sectors (for example in rocketry⁵⁵) due to the conservativeness of this value.

It is important to note however that the number of aircraft operating at any time over the same population at risk will chew through this target proportional to the number of aircraft flying over that population.

7.7 Acceptable levels of safety for third parties in the air

The acceptable level of risk in the air is based upon the current level of Mid-Air Collisions (MAC) over a given period of time. The acceptable level is split between

⁵⁴ Working Group 6 - Safety & Risk Assessment, “AMC RPAS.1309 - Issue 2 Safety Assessment of Remotely Piloted Aircraft Systems.”

⁵⁵ As described in the Flight Safety Code (2018 Department of Innovation, Space and Science (Aust)), the target level of safety for **individual risk** is a probability of death less than 1×10^{-6} per launch.
<https://www.industry.gov.au/sites/default/files/2019-08/space-flight-safety-code.pdf>

controlled and uncontrolled airspace. There are a few reasons for this, namely the larger number of passengers onboard the kinds of aircraft operating in controlled airspace and the higher expectation of safety for regular public transport aircraft (i.e. operations under Part 121 of CASR).

7.7.1 Acceptable levels of safety for third parties in the air within uncontrolled airspace

During 2013, the US Federal Aviation Administration held See and Avoid Workshops to determine the acceptable levels of safety for aircraft in uncontrolled airspace. The second caucus report detailed the requirement for the safety of non-cooperative aircraft (i.e. those most likely to be operating in uncontrolled airspace) as less than 1×10^{-7} MACs per flight hour.⁵⁶

7.7.2 Acceptable levels of safety for third parties in the air within controlled airspace

In controlled airspace, there is a general consensus that the target level of safety adopted for the North Atlantic Tracks (NAT) for Reduced Vertical Separation Minima (RVSM) should apply in controlled airspace. This value is set at 5×10^{-9} fatal accidents per hour for a vertical, longitudinal, or lateral collision risk.⁵⁷ JARUS Draft Annex G sets an acceptable level for MACs in controlled airspace to less than 1×10^{-9} MACs per flight hour.

⁵⁶ FAA, "Sense and Avoid (SAA) for Unmanned Aircraft Systems (UAS), SAA Workshop Second Caucus Report."

⁵⁷ Note that technical systems may need to be assessed against a TLS at half this value. NAT SPG Handbook
https://www.icao.int/EURNAT/EUR%20and%20NAT%20Documents/NAT%20Documents/NAT%20Documents/NAT%20Doc%20001%20-%20NAT%20SPG%20Handbook/NAT%20Doc%20001_NAT%20SPG%20Handbook%20v2.5.0.pdf

7.8 Air Risk and Ground Risk as a function of the Concept of Operations

From a risk perspective, the key difference between personned and unpersonned operations is that in a situation where the aircraft loses control, there are no persons on board that need to be protected from that loss of control. This leads to a phenomenon where the risk is mediated by a combination of the density of third persons in the region of the operation and the damage caused by a crash in the air (the number of people that die in a mid-air collision) or on the ground (the size of the crash area on the ground).

7.8.1 Risk to persons on the ground from UAS operations

Using our pervious equation for risk, we can specify a generic equation that details the i^{th} individual ground risk due to a UAS as:⁵⁸

$$IR_{i, \text{ground risk}} = P(\text{UAS out of control}) \times P(\text{impacts } i^{\text{th}} \text{ person}) \times P(\text{fatally injures } i^{\text{th}} \text{ person})$$

Where:

Term	Explanation
$P(\text{UAS out of control})$	This is the probability that the UAS is out of control and will then impact the ground. This can be affected by a number of factors: • Failure of the UAS in some capacity • Operational error or failure • Failure of a supporting system or service (i.e. communications service provider loss of service) • Adverse operating environment (aircraft operating outside of intended operational conditions)
$P(\text{impacts } i^{\text{th}} \text{ person})$	After the aircraft experiences a loss of control event (row above), in order to fatally injure a person, the aircraft needs to hit someone. This is a function of: • The size of the critical area of the aircraft (i.e. the crash area, the area the UAS covers from the height of a person above the ground until it fully comes to a stop) • The probability that the critical area ends up in a region where it will impact the i^{th} person. This tends to be a function of the population density of the area, as well as the potential area in which the aircraft could crash.
$P(\text{fatally injures person})$	Given that the aircraft has now hit someone, it now needs to fatally injure them (in casualty expectation, this would be injuring to the point of casualty). A conservative starting point is to assume that if the UAS impacts a person, it will kill them (and for larger UAS this is likely a reasonable assumption).


⁵⁸ Note that the equation here contains conditional probabilities, but to keep the equation manageable from a formatting perspective, these have been excluded.

We can sum this individual risk over an entire population at risk, creating the expected fatality equation:

$$E_{F,ground\ risk} = \sum_{\forall i} IR_{i,ground\ risk}$$

Upon some simplifying assumptions⁵⁹, this can be restated for a uniform population density as:

Figure 18 - Ground Risk Expected Casualty Equation




Where λ_{OOC} is the loss of control rate of the operation, or operational reliability (i.e. the expected number of failures per hour), A_C is the critical area of the UAS (i.e. the UAS may have a 5m wingspan and slide during a crash for 20m, resulting in a critical area of 100m²) and D_{pop} is the population density overflowed by the UAS operation. In order to meet the target level of safety (TLOS) for ground risk as described in Section 0 (i.e. less than 1 fatality per million hours of operation), this equation must result in a number lower than the target, i.e.

$$E_{F,ground\ risk} \leq TLOS_{ground\ risk}$$

This means we can combine the loss of control rate, critical area and population density in different ways and still meet the TLOS. For example, we may operate a large aircraft with poor operational reliability over a significantly reduced population density (a flight test range or in a remote area) such that the target level of safety is met:

Figure 19 - Ground Risk Expected Fatality Equation - High failure rate and critical area over smaller population density



⁵⁹ A more thorough representation of the expected casualty (without simplifying assumptions), can be found in Section 4.3.2 of the Flight Safety Code () <https://www.industry.gov.au/data-and-publications/flight-safety-code>

Alternatively, if we wanted to operate that large aircraft over a more populated area, the reliability of the platform would need to increase (which is the same as the loss of control rate being reduced) commensurately to ensure the TLOS is met.

By mediating these terms so that the target level of safety is met, operations of different aircraft sizes and reliabilities can be allowed over population densities that are appropriate to the risk posed by these aircraft.

A critical use of this analysis in risk assessments is to use the critical area of the UAS, the population density and the TLOS to determine the **minimum required reliability of the operation** needed to meet the TLOS for the given combination of UAS and population density:

$$\lambda_{ooc} \leq \frac{TLOS}{A_c \times D_{pop}}$$


7.8.2 Risk to other aircraft from UAS operations

When dealing with airspace risk (i.e. the chance of a mid-air collision) between two aircraft, we first need to define some terms:

Term	Definition
Ownship	The UAS around which the various volumes are constructed. The perspective of encounters are viewed from the ownship perspective.
Intruder Aircraft	An aircraft that intrudes upon the volumes around the ownship. Generally, intruders can be classified as “cooperative” (i.e. are electronically visible and contactable via radio) or “non-cooperative” (i.e. are not electronically visible and are not required to communicate via or may not carry a radio).
Encounter Volume	A volume that is a mandated range for the purposes of consistent analysis for encounter rate (i.e. if the range was allowed to be varied, it would result in different values of the probabilities of MAC, making analysis more complex and convoluted). Within this range mitigations can be applied to reduce the chance of violating the remain well clear volume (RWC volume), i.e. ensuring remain well clear (utilising a RWC function), by self-separation, ATC separation or other means to identify and avoid interacting with intruders within the RWC volume.
Remain Well Clear Volume	This volume defines the boundary between being well clear and not. Within this volume, collision avoidance measures should be taken to prevent an NMAC from occurring (although nothing prevents an aircraft from taking avoidance action at an earlier stage).
Near Mid Air Collision (NMAC) Volume	The NMAC volume defines the boundary within which the probability of a mid air collision is driven by providence.


We can visualise the distance between the ownship and any intruder aircraft as a series of “pucks” or buffers that exists around the ownship (the UAS). As the distance between the intruder aircraft (the conventionally piloted aircraft) and the ownship reduces, more of these buffers will have been violated and the chance of a collision increases:

Figure 20 - An air risk model (equivalent to JARUS Annex D/G model)



We can mathematically represent this as the expected number of mid-air collisions (E_{MAC}) per flight hour (MAC rate or λ_{MAC}):

Figure 21 - The MAC rate equation



Where λ_{enc} is the encounter rate (i.e. the number of aircraft that violate the encounter volume per hour), $P(WCV|enc)$ is the conditional probability that the intruder violates the remain well clear boundary (a **Well-Clear Violation** or WCV) on the condition of experiencing an encounter, $P(NMAC|WCV, enc)$ is the conditional probability that the intruder aircraft and the UAS then experience a Near Mid Air Collision and $P(MAC|NMAC, WCV, enc)$ is the probability that the intruder and UAS, now experiencing an NMAC, collide with each other in a mid-air collision (MAC). The values for all these variables should result in the expected Mid Air Collision rate per hour being below the target level safety.

In a similar manner to the ground risk, aircraft of different types and equipage can mediate the airspace TLOS by operating in areas of lower encounters (i.e. operations below the height of buildings) or by increasing the capability of their equipage to

increase the ability to deconflict with traffic that they encounter (i.e. equipping a detect and avoid system).

Note that this is just one means of dividing up the events before a collision between two aircraft and that other delineations could also be utilised, i.e. introducing a safe separation volume in line with ATC separation requirements or a detection ranges in lines with the equipment of the UAS.


7.9 Bow-tie analysis to visualise hazardous events for UAS

We now have some concept of acceptable levels of safety and the ways in which risk is calculated, how they can be applied to both air and ground risk, and the ways in which we can mediate the risk by managing the who, what, how, and where of the operation in question. These equations we provided above construct these issues simply, however there are complex causal chains that lead to the events we are trying to avoid. In both the ground risk and air risk models, there exist a sequence of events that may end up resulting in an unwanted event (a ground fatality or a mid-air collision). In order to make some form of useful analysis, it is helpful to base any analysis of some kind of differentiator event (which we call the “Top-Level Event”) that can be used to differentiate risk controls and mitigations that are applied before this event and after this event. We could use the “UAS out of control” which would be any state in which the UAS is outside of the intended nominal operating conditions. We can categorise the following types of events that could lead to this top level event as:

- A technical issue with the UAS, or
- A technical issue with an external system that supports the UAS operation (i.e. a loss of communications infrastructure, GPS denial)
- Human error (either part of the UAS operation or otherwise, or
- Adverse environmental conditions (i.e. the UAS operation has exceeded the expected operational limits assumed during the design of the aircraft and operational procedures)
- Another aircraft is on a collision course with the UAS.

Which is usually visualised on a diagram called a fault tree:


Figure 22 - Example of a fault tree that leads to a UAS being out of control



We can also visualise the events that occur after this event in a diagram called an event tree, which depicts a series of binary, casually chained events leading to the following realised outcomes on our desired safety targets:

- Third parties in the air
- Third parties on the ground
- Critical infrastructure


Figure 23 - Event Tree after "UAS out of control" (example only)



We can combine both of these tools into a Bow-Tie Diagram. Bow tie diagrams are a good tool for visualising both sides of the top-level event. On the left of the top-level event are threat lines, these are the things that occur before the top-level event that can lead to the event occurring (and are composed of our fault tree elements). On the right of the top-level event are harms that are likely consequences of the event occurring (composed of our event tree, although necessarily simplified), resulting in a "bow-tie" shape. This visualisation allows the user to see causal lines that lead to top-level events and consequences, and then identify and implements barriers to reduce

the likelihood of these occurring.⁶⁰ The JARUS SORA utilises this method to create their “holistic risk model” for the safety assessment of UAS operations, as shown in Figures 24, 25, and 26.

Figure 24 - SORA V1.0 Bow-tie representation of HRM, p. 16 of SORA V1.0 (JAR-DEL-WG6-D.04)



The top-level event “UAS operation is out of control” is equivalent to λ_{LOC} . The threat lines; technical issue with UAS, deterioration of external systems, human error, aircraft on collision course, and adverse operating conditions all contribute to the top-level event. These threat lines can be driven backwards up the causal chain and threat barriers implemented (i.e. training of human operators to prevent human error) to reduce the likelihood of the top level event occurring. The full SORA bow-tie contains 24 threat barriers, shown below:

⁶⁰ Note that although this is a great tool to visualise risk and determine implementable barriers, it does not allow for the same granularity of quantitative analysis that can be undertaken in a full fault tree or event tree analysis.

Figure 25 - SORA Bow-Tie Model: Threat Barriers





Figure 26 - SORA Bow-Tie Model: Harm Barriers



After this top-level event, the causal chain evolves towards the consequences identified. However, there are still things that can be implemented to prevent the final consequences occurring (i.e. frangibility in the design to prevent an aircraft from killing a person upon impact, parachutes to reduce the crash area and impact energies). These are called harm barriers (the SORA has 4 harm barriers, 3 for ground risk and 1 for air risk). In some bow-tie diagrams, the paths from the top-level event to the outcomes are represented by an event tree.

7.10 The system safety process as a model for safety assurance of technical systems

Two key failure pathways that are identified in the previous sections are technical system based (either the UAS or external systems). We still need a means or method to analyse the constituent parts and processes that lead to the emergence of those top-level failure pathways. In aviation, system safety is the process used to do this.

The system safety process is an engineering process that is aimed at providing safety assurance of systems in both intended and unintended operational environments by providing a means to derive safety objectives and then demonstrate compliance against these objectives.⁶¹ Any system safety process will include the following steps:

- **Hazard Identification** – a process to ensure that all relevant hazards are identified and documented, such that the rest of the system safety process accurately reduces all the hazards associated with the operation.
- **Hazard Analysis and risk mitigation** – after identification of the hazards has been completed, all hazards must be analysed for their consequences and risk mitigations that can be applied.
- **Allocation of safety objectives** – based upon the hazard analysis and the application of risk mitigations, allowable probabilities of failure and software development assurance levels are assigned as the required safety objectives to be met by the safety analysis in order for the system to be considered adequately safe.
- **Verification and validation of safety objectives** – once safety objectives have been set, verification and validation is a process to ensure demonstration of compliance against the objectives by ensuring that the system is built right (verification) and that the right system is built (validation).
- **Continued operational safety** – the safety process does not stop after the design phase. Throughout the operation of the system, review of defects and occurrences to ensure that the system continues to be safe is part of the continued operational safety process.

Whilst this process is applicable to any system in aviation at any risk level, all the current aviation guidance and requirements (for example ARP 4761, ARP 4754) are driven by the risks within conventionally piloted aviation. No consideration is currently given to UAS at different risk levels. Despite this, there are still considerable resources

⁶¹ Note that in conventionally piloted aviation safety, system safety is only a part of the total airworthiness picture, focusing on systems, equipment, and installations only.

to further understand and utilise/modify system safety processes for use in UAS operations:

- [MIL-HDBK-882E](#) – System Safety
- [FAA System Safety Handbook](#)
- [NASA System Safety Handbook](#)
- [FAA AC 23.1309-1E](#) – System Safety Analysis and Assessment for Part 23 Airplanes
- [FAA AC 25.1309-1A](#) – System Design and Analysis
- [SAE Aviation Recommended Practice 4761](#) – Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment
- [SAE Aviation Recommended Practice 4754A](#) – Guidelines for the Development of Civil Aircraft and Systems

7.11 The SORA Risk Assessment, the SAIL and the Aviation Safety Continuum

For unmanned aircraft, high-level, outcome-based requirements without stipulated means of compliance through something like the system safety process makes it difficult to obtain approval to operate these aircraft. This is not an Australia-specific problem and the international community has been trying to grapple with this issue for decades.

A key requirement for any risk assessment process is to contextualise to the proposed operation and all relevant hazards and mitigations that are able to be identified. The JARUS SORA methodology⁶² is a new, standardised risk assessment tool for unpersonned aircraft operations, tailored to the target levels of safety identified in this discussion. Crucially, the JARUS SORA identifies the depth of analysis required for all the barriers that could be implemented in a UAS operation, based upon the system safety concept, providing a tool for truly risk-based safety regulation for UAS.

CASA has adopted the use of the JARUS SORA as for use to inform its policies on complex RPAS operations.⁶³ Following a robust SORA assessment will allow the highest probability of complying with the CASA approval requirements for any complex highly automated or autonomous operation.

The SORA follows the following steps, which utilises methodologies similar to what has been presented in this document throughout Section 10.6:

⁶² JARUS Specific Operations Risk Assessment v2.0 and associated annexes can be found here: jarus-rpas.org/content/jar-doc-06-sora-package

⁶³ CASA, “Safe Integration of RPAS into Australian Skies.”

Concept of Operations Description

Step #1. ConOps Description

Determination of risk posed to third parties on the ground (See 7.6 and 7.8.1)

Step #2. Determination of the Intrinsic Ground Risk Class (GRC)

Step #3. Final GRC determination

Determination of risk posed to third parties in the air (See 7.7 and 0)

Step #4. Determination of the Initial Air Risk (ARC)

Step #5. Applications of Strategic Mitigations

Step #6. Tactical Mitigation Performance Requirements (TMPR)

Determination of overall risk level and allocation of acceptable depth of analyses requirements (See, □, 0 and 0)

Step #7. SAIL Determination


Step #8. Identification of Operational Safety Objectives

Ensuring that risk is acceptable to those outside of the designated approved area

- Adjacent Area/Airspace Considerations

Steps #2, #3, #4, #5 and #6 in the SORA process follow similar air and ground risk models discussed in this document to effectively “place” the operation along the aviation safety continuum into categories called Specific Integrity and Assurance Levels (SAILs), assigned in Step #7, based upon the minimum required reliability needed by the operation (similar to λ_{00C} described in 7.8.1). This is shown below:

Figure 27 - SORA SAILs shown on Safety Continuum⁶⁴



⁶⁴ Note that the position of an operation that kills at least one person per catastrophic event between SAIL IV and SAIL V does not imply that this is always the case (it depends on a range of factors beyond the scope of this analysis).

The SORA uses the bow-tie model described in 7.9 to allocate 24 assurance requirements (called operational safety objectives or OSOs) to the operation, across all aspects of the (technical, operational, organisational etc.) based upon the level of safety required by the SAIL level. The key characteristic of the SORA that makes it a true risk-based approach is the use of **robustness levels** to dictate depth of analysis needed to demonstrate compliance with the safety objectives.

At each SAIL level, the OSOs must be met at the required robustness levels as specified by Table 6 in SORA V2.0.⁶⁵ Robustness is defined as a combination of:

- The amount of safety gain or benefit a system gives (the level of integrity); and
- The certainty that the integrity claim is true (the level of assurance).

And are categorised into three levels of robustness:

- **Low Robustness** – for lower risk activities, low robustness relies upon self-declaration against a set of integrity requirements.
- **Medium Robustness** – medium robustness required evidence to be shown to a regulator that a given set of integrity requirements have been met. This may require compliance with standards but may not require third party validation of claims against the standard.
- **High Robustness** – high robustness will require competent third-party verification against a recognised standard. At high robustness, the standard will approach or be equivalent to conventionally piloted aircraft standards.

The 24 assurance requirements (called Operational Safety Objectives, or OSOs in SORA parlance) are expert identified threat barriers specific to UAS operations and are related to the five threat lines. As the risk increases, and the operation moves along the risk continuum, the required objectives increase in number and required robustness level. The table below shows the robustness of all 24 OSOs as a function of the SAIL level.

Table 3 - SORA v2.0 SAIL Level Objective Requirements

SAIL Level	I	II	III	IV	V	VI
Optional Objectives	9	6	1	0	0	0
Low Robustness Objectives	15	14	6	1	0	0
Medium Robustness Objectives	0	4	13	16	4	0
High Robustness Objectives	0	0	4	7	20	24

⁶⁵ JARUS, "JARUS Guidelines on Specific Operations Risk Assessment (SORA)."

In this way the SORA provides a complete method to measure aviation risk, determine requirements and specify the depth of analysis required to meet those identified requirements.

The SORA is an extensive body of work created by many international participants, and significant bodies of work exist on this topic. Readers should access the JARUS [publications webpage](#) for more information, guidance on its use, and more.

8 Assurance issues for highly automated and autonomous UAS

Section 8 highlights the assurance issues facing highly automated and autonomous systems that need to be solved to truly allow the integration of these operations into Australian airspace.


Even with the SORA process providing for a proportionate risk assessment framework for UAS, there are still gaps and issues for current operations and future autonomous operations. These need to be addressed to fully facilitate the safety assurance of highly automated and autonomous systems.

As described, the scope of autonomy is extremely large and covers many facets. Referring back to the concepts laid out by Clothier, Williams and Perez (2019), we can try to understand the challenges facing the assurance of highly automated and autonomous systems using the dimensions analysed by this paper and leading into the expansion of operations into the future.

As a reminder, the current state of autonomy for UAS is restricted to the realms of:

- Deterministic (Model Based, Specified Algorithms) – i.e. at some level a human has specified an algorithm detailing how these systems should work; and
- The human as the primary authority – The final decisions for any of these systems is determined by a human, either during operation, or during the design process.

Figure 28 - Current extent of autonomy of UAS operations



Resulting in most operations restricted to the information analysis (sense, perceive) and action levels of functional performance (i.e. any information analysis or decision/action selection is restricted to pre-programmed responses).

From a risk perspective, these operations are still extremely limited in operational scope for use in air and on ground, i.e. under SORA terminology, these operations are usually restricted to SAIL I and SAIL II operations.

As current best practice for assurance for systems utilises system safety approaches, we will categorise the challenges for assurance aligned with this process. Note that

this list is not exhaustive and there are likely more issues that need to be addressed to fully capture the risks and requirements appropriate to highly automated and autonomous UAS.

8.1 Incorporating analysis of autonomy into the SORA and system safety process

The system safety process is a holistic, well -understood and -utilised framework for engineering risk analyses of systems (i.e. MIL-STD-882E, ARP 4761, IEC 61508, ISO 26262) that can be applied to highly automated/autonomous UAS. However, the current system safety process in aviation is focused (understandably) on the risks in conventionally piloted aviation. The SORA process, introduced in 7.11, is an attempt to apply a system safety methodology to the entirety of UAS operations, including operational and organisational risks. Due to this, there still exist gaps and issues with respect to the incorporation of highly automated/autonomous systems into this process. In order to highlight and specify these gaps we will use the stages in the system safety process highlighted in 7.10.

8.1.1 Hazard Identification (HAZID), Hazard Analysis and Risk Mitigation

Hazard Identification (HAZID) is a process to ensure that all relevant hazards are identified and documented, such that the rest of the system safety process accurately reduces all the hazards associated with the operation.

HAZID processes usually rely upon checklists or expertise garnered through operations and failures. The expansion of the system functions to those currently undertaken by persons means that the previously acceptable checklists or expertise may miss hazards that are introduced by the use of autonomous systems (i.e. cognitive errors or lapses in information analysis or decision making processes).

After the HAZID process, all hazards must be analysed for their consequence, and risk mitigations identified for inclusion into the CONOPS. Hazard analysis is predicated on an understanding of the outcomes of hazards (when they occur) and aligning safety expectations based upon these outcomes.

To correctly identify and analyse the consequences of new and arising hazards as a result of increasing autonomy in aviation there needs to be additional guidance, standards, and processes to ensure that all credible hazards specific to higher autonomy systems are captured in the safety assessment.

When it comes to risk mitigation, new autonomous systems – with novel components and architectures – will need a raft of unique mitigations. Guidance on useful mitigations for autonomous system, including how to ensure the success of these mitigations is needed. Some examples could include:

- Bounding a machine learning system with traditional, certified, model-based systems that only allow outputs within a pre-determined “safe” range.
- Bounding a complex-model based algorithm (i.e. a path-finding algorithm) with traditional, certified, model-based systems that engages when specific conditions are reached.

- Modular design, redundancy, separation, independency, specific to autonomous systems (i.e. using two or three independent algorithms for the same purpose, comparing outcomes and using the consensus/fused result)

8.1.2 Allocation of safety objectives

Based upon the hazard analysis, including the application of risk mitigations, allowable probabilities of failure and software development assurance levels are assigned as the required safety objectives that need to be met to ensure an acceptable level of safety. These objectives are driven by the consequence of the top-level event occurring. As we have discussed, if we want a true risk based and proportionate framework, these objectives need to be tailored for UAS operations along a spectrum. A catastrophic event in one operation does not mean a catastrophic event in another (i.e. hull loss for a large transport category aircraft could kill hundreds of people, whilst the hull loss of a small UAS over rural areas may not kill anyone). The ability to scale the allocation of safety objectives, particularly for lower safety risk is needed.

It is important to note that many assumptions go into the allocated safety objectives. As an example, we can analyse the catastrophic objective for a Part 23 (normal, acrobatic and commuter category) Class I aircraft. This objective states any catastrophic failure should not occur more than one in every million flight hours:

$$\lambda_{CAT, Pt\ 23\ Class\ I} \leq 1 \times 10^{-6} pfh$$

This is built up from the following analysis:⁶⁶

1. The current fatal accident rate experienced by single engine, <6,000lb aeroplanes is 1 in 10,000 flight hours (1×10^{-4} per flight hour):

$$\lambda_{fatal\ accident} = 1 \times 10^{-4} pfh$$

2. 90% of these events are caused by human error (HE), 10% by the machine system (MS) . We can split the top-level fatal accident rate into each type of failure (human or machine):

$$\lambda_{fatal\ accident} = 10 \times \lambda_{fatal\ accident, MS}$$

$$\lambda_{fatal\ accident, MS} \leq 1 \times 10^{-5} pfh$$

3. There are 10 machine system failure modes (each a catastrophic event) that could individually lead to these fatal accidents, and therefore each of these individual events must sum to less than this number. An assumption is made

⁶⁶ FAA, "AC 23.1309-1E - System Safety Analysis and Assessment for Part 23 Airplanes – Document Information."

that each of these take up an equal portion of the allowable machine system contribution to the fatal accident rate:

$$\lambda_{fatal\ accident,MS} \ni \{\lambda_{CAT1}, \lambda_{CAT2} \dots \lambda_{CAT10}\}$$

$$\forall i, \lambda_{CATi} = \lambda_{CAT}$$

$$\lambda_{fatal\ accident,MS} = 10 \times \lambda_{CAT}$$

- And, therefore, we arrive back at the original requirement:

$$\lambda_{CAT,Pt\ 23\ Class\ I} \leq 1 \times 10^{-6} pfh$$

These assumptions may not hold when we look at increasingly autonomous operations:

- **The contribution of human error to the fatal accident rate:** As the level of autonomy increases, the contribution of human error to the accident rate will likely change. There is currently no concrete guidance on what this means. One solution is to leave the requirement for machine systems at the level derived earlier,⁶⁷ other solutions may allow the machine system to “absorb” a portion the allowable failure rate that is no longer incurred through human error.
- **The number of machine system failure modes that could lead to a fatal accident:** As systems become more complex, there is the possibility that the number of potential hazards that could lead to a fatal accident could significantly exceed the assumed 10 failure conditions from AC 23.1309-1E. There needs to be a way to incorporate the number of failure modes into the assessment (potentially through the functional hazard analysis process) such that the target levels of safety are achieved.

Additionally, in the 1970’s it was discovered that complex hardware and software have a large risk of incurring development errors that lead to failures. These are not readily analysable with traditional reliability techniques and development assurance, which were included as part of software development. Learning based algorithms are exposed to similar risks (i.e. development errors), but require different techniques to ensure that adequate protection from these risks is included in the development process. A scalable learning assurance process is needed (see 0), proportionately allocated as a function of the consequence of an event.

⁶⁷ Working Group 6 - Safety & Risk Assessment, “AMC RPAS.1309 - Issue 2 Safety Assessment of Remotely Piloted Aircraft Systems.”

8.1.3 Verification and validation of safety objectives

The previous stages in the system safety process were all about deriving safety objectives based upon the hazards that an operation experiences. Once safety objectives have been set, verification and validation processes are then used to demonstrate compliance against the derived objectives. Verification is the process to ensure that the system is built right (the design meets the objectives) and validation is the process used to ensure that the right system was built (the safety objectives are still met by the designed system). We identify challenges in the current process for highly automated/autonomous systems below (and by no means is this list exhaustive!):

8.1.4 Model-based software assurance for highly automated/autonomous UAS

Even within traditional, model-based software assurance, there are issues with assurance of highly automated or autonomous systems. These include:

- Development assurance levels and standards specific to UAS

The current development assurance level attribution is based off crewed aircraft operations. As discussed, the current allocation of safety objectives do not scale for lower risk UAS operations, leading to a strange limbo state where lower risk UAS operations need to either:

1. Comply with an overly burdensome development assurance process (i.e. DAL C under RTCA DO-178C) that does not take into account the lower risk of the UAS operation, or
2. Propose or undertake software development assurance to a standard that is not accepted and attempt to make a safety case justification to the regulator based upon the use of that standard

Neither of these solutions are particularly helpful to a designer and a list of approved development assurance standards for risk levels in UAS that are not covered by current standards is needed.

- Verifying software functionality for complex task in complex task-environments

When the task or task-environment increases significantly in complexity, the ability to demonstrate that a piece of model-based software functions as intended across the entire possible task environment can be incredibly difficult, if not impossible due to issues such as the combinatorial explosion⁶⁸ or satisfiability problems.⁶⁹

We need methods to safely bound the performance of these complex model-based functions, at a level commensurate to the risk of the operation. Guidance on incorporating analytical techniques (formal methods – formal specification, development, and verification) and simulation/test techniques (i.e. monte-carlo, boundary, or edge case testing) are needed to assist in the demonstration of software performance across the task environments that highly automated/autonomous systems will be designed for.


⁶⁸ Wikipedia, "Combinatorial Explosion."

⁶⁹ Wikipedia, "Boolean Satisfiability Problem."

8.1.5 Moving to data-based (pre-trained and adaptive) algorithms for safety critical functions


Data-based algorithms are fundamentally different from traditional, model-based algorithms that are used in safety critical functions in aviation today. We can think of a system as an agent that responds to inputs (I) by applying some function (f) that maps the inputs to outputs (O):

Figure 29 - A model for how an autonomous system function works




For all inputs i in the possible input space there exists an output o , in all possible output space, such that the mapping function f applied to the input value i results in output o . In model-based algorithms, the human designer directly and explicitly defines the function that creates the behaviour of the system.

Figure 30 - model-based systems



In data-based algorithms, a learning algorithm generates this function instead, based upon a set of training data. The human instead may explicitly define the learning algorithm and the dataset (although in some reward modelling scenarios, another learning algorithm may approximate these too) that is used creates the function:

Figure 31 - data-based or learning systems



Traditional software assurance processes assume that the type of implementation is model-based, meaning that a significant amount of development assurance processes are not appropriate or useful for assuring data-based systems. A “learning assurance” process is needed, specific to data-based systems, which provides a risk proportionate process for developing these systems. Some things of note that need to be covered include (this list is not exhaustive):

- **Data management:** Learning systems live and die off their training, so the datasets used for the test, training and validation of the learning process must adequately cover the total feature space of the operational environment and must be constructed in such a way that each set is orthogonal to one another
- **Performance measures for learning systems:** Performance measures – such as generalizability, bias, and variance – can give insight into the performance of the learning algorithm and its ability to represent an underlying, unknown model.⁷⁰ Importantly, we want these systems to be able to create a function that works acceptably well across the entire operational environment and is not brittle to small changes in inputs from trained datapoints.
- **The alignment problems:** Misalignment between the true, intended goals of the human designer and the actual implemented goals in the training process (outer alignment⁷¹) and the misalignment between the implemented goals during training and the actual goals of the learning agent (inner alignment).⁷²
- **Explainability/Transparency/interpretability:** The function created by the learning system can be quite opaque to the human designer (i.e. in the worst cases it can be a true black box), so there needs to be a way to understand the reasons why certain decision or outputs are achieved by the learned function

⁷⁰ Cluzeau et al., “Public Report Extract Concepts of Design Assurance for Neural Networks (CoDANN).”

⁷¹ Clark and Amodei, “Faulty Reward Functions in the Wild.”

⁷² Heaven, “Why Deep-Learning AIs Are so Easy to Fool.”

in order for analysis and interpretation to occur. **Why** something is chosen or determined is an important question, the answer to which (if accurate!) engenders trust in the decision-making process.

- **Pre-Trained vs Adaptive Learning:** A key differentiator within the realms of learning systems are those that are pre-trained (i.e. they are trained during specific non-operational epochs, and frozen in time during operation) and those that are adaptive (i.e. they continually learn and adapt their function from inputs during operation). Current certification processes are predicated upon the concept of approving a design that is frozen at the point of certification. For pre-trained learning systems, this is not a problem, however for adaptive learning systems, this poses a significant challenge and an understanding of how adaptive systems can be approved is needed.

8.1.6 Digital Twins and Digital Environment Verification and Validation

The increase in task-environment complexity that highly automated/autonomous systems must be able to operate within, alongside the increased use of complex model-based algorithms or learning algorithms will require advanced simulation of:

- The system or agent (including the ability to sense the environment through sensors, replicating the agent function, and act on the environment with actuators);
- The task-environment (i.e. all external stimuli and entities/components that can interact with or are modified by the agent); and
- The interaction between the system and the task-environment.

To ensure that any inferences made from the simulation represent the true inferences that can be made in reality, whether for test, demonstration or learning using synthetic datasets, guidance, processes and standards are needed to ensure adequate replication of reality.

8.1.7 Common Cause Analysis for new applications of systems

In any functional hazard analysis, there is a huge assumption on independency of events. This assumption is currently managed by the use of common cause analysis (CCA) techniques:

- Zonal hazard analysis;
- Particular risk analysis; and
- Common failure mode analysis.

These are all predicated on analysing the assumption of independence and ensuring that mitigations are put in place to retain this assumption. The increase in cognitive task complexity may introduce new interdependencies that cannot be identified by current CCA techniques. An analysis of or introduction of a means to identify dependent events in complex tasks/task-environments is critical to reducing the risk of common cause failure modes in a future highly automated/autonomous system safety process.

9 References

- AIHW. "Deaths in Australia." Australian Institute of Health and Welfare, 2021. <https://www.aihw.gov.au/reports/life-expectancy-death/deaths-in-australia/data>.
- Airservices Australia. "Flight Information Management System (FIMS): Request for Proposal | Engage Airservices," 2021. <https://engage.airservicesaustralia.com/fims-rfp>.
- Assure UAS. "Assure - THE FAA's Center of Excellence for UAS Research," 2022. <https://www.assureuas.org/>.
- CASA. "CASA's Regulatory Framework." DCMIType. Civil Aviation Safety Authority. Civil Aviation Safety Authority, 2021. <https://www.casa.gov.au/rules/regulatory-framework/casas-regulatory-framework>.
- . "Our Regulatory Philosophy," 2022. <https://www.casa.gov.au/about-us/who-we-are/our-regulatory-philosophy>.
- . "Safe Integration of RPAS into Australian Skies." DCMIType. Civil Aviation Safety Authority. Civil Aviation Safety Authority, June 21, 2021. <https://www.casa.gov.au/about-us/content-search/news-media-releases-speeches/safe-integration-rpas-australian-skies>.
- Clark, J, and D Amodei. "Faulty Reward Functions in the Wild." OpenAI, December 22, 2016. <https://openai.com/blog/faulty-reward-functions/>.
- Cluzeau, J. M., X Henriquel, G Rebender, and G Soudain. "Public Report Extract Concepts of Design Assurance for Neural Networks (CoDANN)," 2020. <https://www.easa.europa.eu/sites/default/files/dfu/EASA-DDLN-Concepts-of-Design-Assurance-for-Neural-Networks-CoDANN.pdf>.
- Commonwealth of Australia. "Work Health and Safety Act 2011." Attorney-General's Department, 2011. Au. <https://www.legislation.gov.au/Details/C2018C00293/Html/Text>, <http://www.legislation.gov.au/Details/C2018C00293>.
- Cour-Harbo, A Ia. "Mass Threshold for 'Harmless' Drones -." *International Journal of Micro Air Vehicles*, 2017, 77–92. <https://doi.org/10.1177/1756829317691991>.
- DASA. "Risk Controls for UAS Operations." Department of Defence, 2018. <https://defence.gov.au/DASP/Docs/DASR-Documents/AA-Circulars/AdvisoryCircular001-18-RiskControlsforUASOperations.pdf>.
- Department of Defence. "DASR - Aviation Safety Management System: Aviation Safety Management Systems." Website. Defence Aviation Safety Authority. scheme=AGLSTERMS.AglsAgent; corporateName=Commonwealth of Australia, Department of Defence; address=Russell Offices, Russell, ACT, 2600; contact=1300 333 362, 2021. <https://defence.gov.au/DASP/ASMS/AviationSafetyManagementSystem.asp>.
- . "Defence Aviation Safety Regulations (DASR)." Defence Aviation Safety Regulation, 2021. <https://defence.gov.au/DASP/Docs/Manuals/DefenceAviationSafetyRegulation/DASRWeb/index.htm#15949.htm>.
- Department of Infrastructure, Transport, Regional Development and Communications. "National Emerging Aviation Technologies Policy Statement." Commonwealth of Australia, 2021. <https://www.infrastructure.gov.au/sites/default/files/documents/national-emerging-aviation-technologies-policy-statement.pdf>.

- Department of Transportation. Federal Aviation - 14 CFR Parts 21, 43, 61, et al. Operation and Certification of Small Unmanned Aircraft Systems, Vol 81 No 124 § (2016). <https://www.govinfo.gov/content/pkg/FR-2016-06-28/pdf/2016-15079.pdf>.
- EASA. "Certified Category - Civil Drones | EASA." European Union Aviation Safety Agency, 2021. <https://www.easa.europa.eu/domains/civil-drones/drones-regulatory-framework-background/certified-category-civil-drones>.
- . "EASA Artificial Intelligence Roadmap 1.0 Published." European Union Aviation Safety Agency, 2020. <https://www.easa.europa.eu/newsroom-and-events/news/easa-artificial-intelligence-roadmap-10-published>.
- . "EASA Concept Paper: First Usable Guidance for Level 1 Machine Learning Applications." European Union Aviation Safety Agency, 2021. https://www.easa.europa.eu/sites/default/files/dfu/easa_concept_paper_first_usable_guidance_for_level_1_machine_learning_applications_-_proposed_issue_01_1.pdf.
- . "EASA Releases for Consultation Its First Usable Guidance for Level 1 Machine Learning Applications." EASA, April 21, 2021. <https://www.easa.europa.eu/newsroom-and-events/news/easa-releases-consultation-its-first-usable-guidance-level-1-machine>.
- . "Open Category - Civil Drones | EASA." European Union Aviation Safety Agency, 2021. <https://www.easa.europa.eu/domains/civil-drones/drones-regulatory-framework-background/open-category-civil-drones>.
- . "Specific Category - Civil Drones." European Union Aviation Safety Agency, 2021. <https://www.easa.europa.eu/domains/civil-drones/drones-regulatory-framework-background/specific-category-civil-drones>.
- FAA. "AC 23.1309-1E - System Safety Analysis and Assessment for Part 23 Airplanes – Document Information." Federal Aviation Administration, 2011. https://www.faa.gov/regulations_policies/advisory_circulars/index.cfm/go/document.information/documentID/1019681.
- . "Certification for Advanced Operations Unmanned Aircraft Systems (UAS) – Criteria for Special Classes." Federal Aviation Administration, 2021. https://www.faa.gov/uas/advanced_operations/certification/criteria_special_classes/.
- . "Sense and Avoid (SAA) for Unmanned Aircraft Systems (UAS), SAA Workshop Second Caucus Report," January 2013.
- Heaven, Douglas. "Why Deep-Learning AIs Are so Easy to Fool." *Nature* 574, no. 7777 (October 9, 2019): 163–66. <https://doi.org/10.1038/d41586-019-03013-5>.
- ICAO. "Convention of International Civil Aviation." ICAO, December 7, 1944. https://www.icao.int/publications/Documents/7300_orig.pdf.
- . "Global Air Traffic Management Operational Concept," 2005. [https://www.icao.int/Meetings/anconf12/Document%20Archive/9854_cons_en\[1\].pdf](https://www.icao.int/Meetings/anconf12/Document%20Archive/9854_cons_en[1].pdf).
- . "Unmanned Aircraft Systems (UAS)," 2011. 328. https://www.icao.int/meetings/uas/documents/circular%20328_en.pdf.
- JARUS. "JARUS Guidelines on Specific Operations Risk Assessment (SORA)," 2019. http://jarus-rpas.org/sites/jarus-rpas.org/files/jar_doc_06_jarus_sora_v2.0.pdf.
- NASA. "UAS Traffic Management (UTM) Project." Text. NASA, 2018. <http://www.nasa.gov/utm>.
- Porter, M, and N Morgenstern. "Landslide Risk Evaluation in Canada." Banff, AB, Canada, 2012.

- https://www.researchgate.net/publication/280445256_Landslide_Risk_Evaluation_in_Canada.
- Range Commanders Council. "Common Risk Criteria Standards for National Test Ranges, Standard 321-20," 2020.
https://www.trmc.osd.mil/wiki/download/attachments/113019899/321-20_Supplement.pdf?version=1&modificationDate=1623181167467&api=v2.
- United Nations. United Nations Convention on the Law of the Seas, Article 3 § (n.d.).
https://www.un.org/depts/los/convention_agreements/texts/unclos/unclos_e.pdf.
- Wikipedia. "Boolean Satisfiability Problem." In *Wikipedia*, January 9, 2022.
https://en.wikipedia.org/w/index.php?title=Boolean_satisfiability_problem&oldid=1064645537.
- . "Combinatorial Explosion." In *Wikipedia*, February 13, 2022.
https://en.wikipedia.org/w/index.php?title=Combinatorial_explosion&oldid=1071678703.
- Working Group 6 - Safety & Risk Assessment. "AMC RPAS.1309 - Issue 2 Safety Assessment of Remotely Piloted Aircraft Systems." JARUS, November 2015.
https://rpas-regulations.com/wp-content/uploads/2017/03/151115_JARUS_AMC-RPAS-1309_issue_2_2.pdf.