



Ministry of Information Technology, Communication and Innovation

FAIR

DIVA
(Digital Interactive Virtual Assistant)
your AI Partner



**GUIDELINES
FOR THE DEVELOPMENT AND USE
OF ARTIFICIAL INTELLIGENCE**



CONTENTS

2 >	2 >	3 >	3 >
Introduction to the FAIR guidelines	Navigating through the Guidelines	Purpose, Scope and Legal Trajectory of the FAIR Guidelines	Who the FAIR Guidelines Are For
> 3 >	8 >	8 >	10 >
Scope of Application	Guiding Values- The FAIR Pillars	Roles and Responsibilities - A Shared Governance Model	AI System Lifecycle - From Design to Decommissioning
> 12 >	13 >	14 >	14 >
Risk-Based Governance Model	Indicative Risk Tiers	Low-risk systems	Medium-risk systems
> 15 >	15 >	17 >	20 >
High-risk systems	Unacceptable risk	Escalation, Controls and Safeguards	Data Governance and Data Protection Alignment
> 22 >	24 >	26 >	28 >
Transparency, Explainability and Human Oversight	Accountability, Auditability and Record-Keeping	Procurement, Third-Party and Vendor Governance	Generative, Foundation and Agentic AI Systems
> 30 >	32 >	33 >	35
Sectoral Adaptation and Implementation Principles	Capacity Building and Competence	Monitoring, Review and Continuous Improvement	From Guidelines to Policy and Legislation

Introduction to the FAIR guidelines

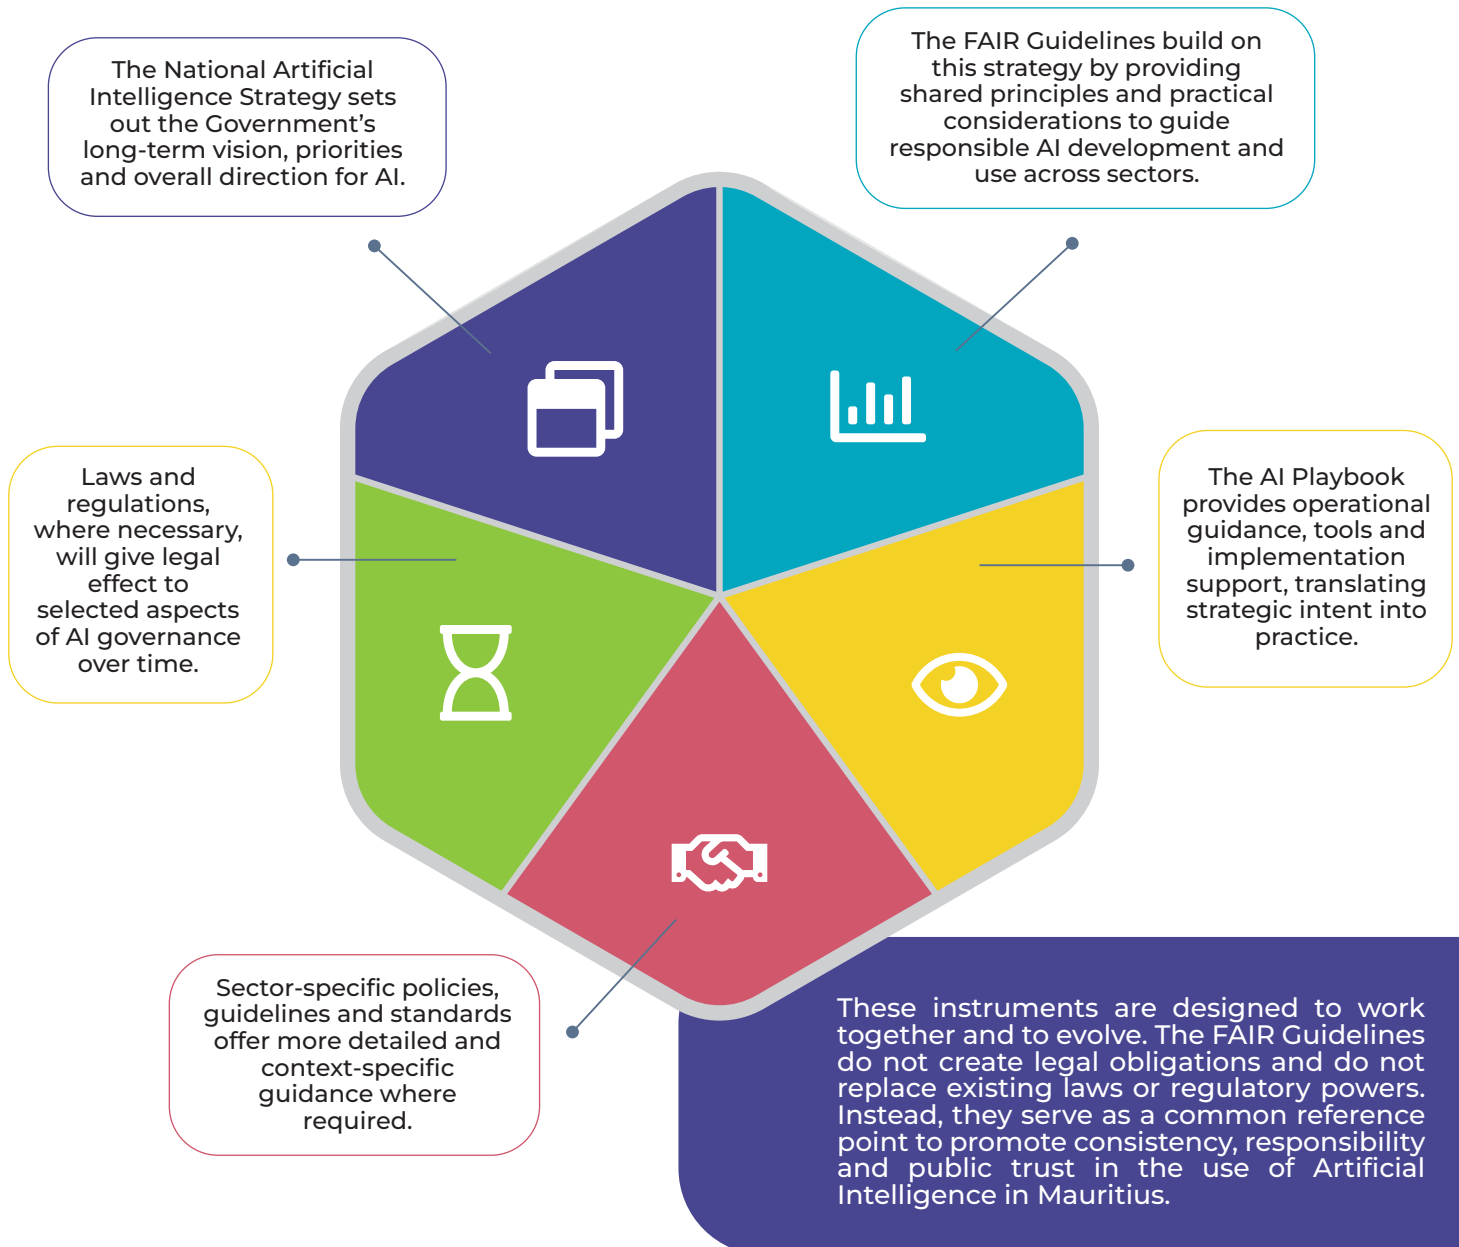
The FAIR Guidelines for the Development and Use of Artificial Intelligence set out a shared, non-binding approach to the responsible and trustworthy development and use of AI in Mauritius.

They are intended to help public institutions, private organisations, technopreneurs and other stakeholders understand what responsible AI use means in practice and how potential risks and impacts should be identified and managed as AI is adopted across the economy and society.

The FAIR Guidelines are not a strategy, a law or an operational manual. They are designed to guide good practice, support informed decision-making and promote public trust, while allowing flexibility across sectors, use cases and levels of risk.

Navigating through the Guidelines

Our approach to Artificial Intelligence is supported by a set of complementary policy instruments that operate at different levels.



Purpose, Scope and Legal Trajectory of the FAIR Guidelines

Purpose of the FAIR Guidelines

The purpose of the FAIR Guidelines is to support a clear, coherent and nationally consistent approach to the development and use of Artificial Intelligence in Mauritius.

As AI technologies are increasingly adopted across public administration, the private sector and society more broadly, there is a need for a shared governance reference that clarifies expectations, responsibilities and safeguards, without constraining innovation or pre-empting future legal or technological developments.

The FAIR Guidelines respond to this need by setting out the principles, governance logic and risk-based approach that should guide how AI systems are designed, developed, deployed and used in Mauritius. They provide a common foundation to support responsible adoption, build public trust and ensure that AI contributes positively to national development objectives.

These Guidelines are intentionally practical rather than theoretical. They reflect the realities of a small, open economy with limited administrative capacity, while remaining aligned with evolving international best practice. They are designed to be used, adapted and built upon by institutions and sectors over time.

The FAIR Guidelines serve as a common reference to support responsible and trustworthy development and use of Artificial Intelligence across different sectors and stages of the AI lifecycle.

Who the FAIR Guidelines Are For

The FAIR Guidelines are relevant to a wide range of actors involved in the development, deployment, procurement or use of AI systems, including:

- **Public sector institutions and state-owned enterprises**
- **Regulators and oversight bodies**
- **Private sector organisations, including micro enterprises, SMEs, start-ups and large firms**
- **Technology developers, vendors and service providers**
- **Academic and research institutions**

Civil society organisations and other stakeholders engaged in AI-related activities

Organisations and sectors using AI are encouraged to draw on the principles and guidance set out in these Guidelines and to adapt them, as appropriate, to their specific context, scale and level of risk.

Scope of Application

The FAIR Guidelines apply across all sectors and to all stages of the AI lifecycle.

They apply to AI systems developed locally as well as those sourced from external providers. Where AI systems are used in Mauritius or where their use affects individuals, organisations or public interests in Mauritius, the principles and governance expectations set out in these Guidelines are intended to apply.

The Guidelines cover the full lifecycle of an AI system, from early conception and design through deployment, operation, monitoring and, where appropriate, decommissioning. Governance expectations are intended to scale with context, impact and risk, recognising that not all AI systems require the same level of oversight.

The FAIR Pillars

The **FAIR Guidelines** are grounded in four core values:



Together, these pillars provide a practical foundation for the responsible development and use of Artificial Intelligence in Mauritius.

The FAIR pillars are intended to guide real decisions, not to serve as abstract principles. They shape how AI systems should be designed, developed, deployed and overseen across sectors and how risks and impacts should be managed in practice. They apply throughout the AI lifecycle and should be interpreted in a manner proportionate to context, scale and risk.

The pillars reflect national values, constitutional principles and public service traditions, while remaining consistent with evolving international practice. They are designed to be technology-neutral and capable of informing future policy and legal measures as our AI ecosystem matures.

Fairness

Fairness requires that AI systems are developed and used in a manner that avoids unjustified bias, discrimination or exclusion and that respects equality, proportionality and legitimate expectations.

In practice, fairness recognises that AI systems can reinforce or amplify existing inequalities if risks are not identified and addressed. Organisations are therefore expected to consider how data, models and automated outputs may affect different individuals or groups, particularly where decisions have meaningful or lasting consequences.

Fairness does not require identical treatment in all cases. It requires informed judgement, taking into account the purpose of the AI system, the context in which it is used and the potential impact of errors or bias. As impact and risk increase, expectations around testing, monitoring and mitigation also increase.

Accountability

Accountability requires that responsibility for AI systems is clearly defined, traceable and retained throughout their lifecycle. AI systems should not obscure who is responsible for decisions or outcomes. For each system, there should be clarity regarding who is responsible for its development and deployment, who oversees its operation and who remains accountable when issues arise.

Accountability includes appropriate human oversight, ensuring that AI systems support rather than replace human judgment, particularly in higher-risk or high-impact contexts. It also includes appropriate transparency and explainability, proportionate to risk, so that the use of AI systems can be understood and reviewed. Mechanisms for review, challenge and redress should be available, particularly where AI systems influence decisions that affect people or organisations.

This pillar reflects a simple principle: the use of AI does not remove the need for human judgement, institutional ownership or clear lines of responsibility.

Inclusiveness and Integrity

Inclusiveness and integrity ensure that AI contributes to broad societal benefit while maintaining trust in institutions, markets and public decision-making. Inclusiveness recognises that the benefits of AI should not be limited to a small number of actors. It encourages attention to access, participation and capacity-building, particularly for public institutions with limited resources, small and medium-sized enterprises and groups that may otherwise be excluded.

Integrity requires that AI systems are developed and used in line with the rule of law, data protection requirements and ethical standards. This includes ensuring lawful data practices, system security, robustness and resilience and protection against misuse or unintended harm.

Together, inclusiveness and integrity support confidence in the way AI is introduced and used and help ensure that adoption strengthens trust rather than undermining it.

Responsibility

Responsibility requires that AI systems are developed, deployed and managed with due care for their effects over time.

Responsibility does not end once a system is deployed. Organisations are expected to monitor how AI systems perform in practice, keep appropriate records, support audit and review and respond to unintended outcomes. Where risks become unacceptable, systems should be adjusted, limited or withdrawn.

This pillar also encourages foresight. It calls for consideration of longer-term and cumulative effects, particularly where AI systems are used in essential services or decision-making processes. By treating responsibility as an ongoing obligation rather than a one-off assessment, this pillar supports learning, adaptation and continuous improvement across the AI lifecycle.

How to Use These Guidelines



These Guidelines are intended to be used as a practical reference rather than as a set of prescriptive rules. Readers may find it helpful to approach them in the following way.

Principles first

The FAIR pillars - Fairness, Accountability, Inclusiveness and Integrity and Responsibility - provide the core values that underpin the Guidelines. Each pillar is informed by considerations such as safety, transparency, security, non-discrimination, contestability and redress. These considerations help clarify how the principles should be understood and applied in practice, without creating separate or standalone obligations.

Risk awareness over time

The performance and impact of AI systems may change as data, behaviour and operating conditions evolve. Information that is accurate or appropriate at the time of development may become outdated or less reliable. As the potential impact of an AI system increases, greater care, oversight and periodic review may be appropriate throughout its lifecycle.

Context matters

Different sectors and organisations operate under different conditions and constraints. These Guidelines provide general, cross-cutting guidance rather than one-size-fits-all rules. Regulators and organisations are encouraged to draw on them and adapt their approach in a way that reflects their specific responsibilities, capacities and risk environment.

Evolution over time

AI technologies and uses continue to evolve. These Guidelines are intended to support learning, review and continuous improvement and may be updated over time as experience, evidence and institutional capacity develop.

By maintaining a clear distinction between national strategy, guiding principles and practical guidance, Mauritius seeks to promote a responsible, flexible and forward-looking approach to Artificial Intelligence that remains aligned with national values and international best practice.



Legal Trajectory and Policy Evolution

The FAIR Guidelines are designed with a **clear policy and legal trajectory**.

While the Guidelines themselves are non-binding, it is intended to inform and support the progressive development of:

- Government policies and administrative guidance
- Sector-specific regulatory instruments
- Procurement standards and contractual requirements
- Future legislative measures, where justified by risk, scale or societal impact

By articulating governance principles and expectations at this stage, the Guidelines create a stable reference point that can be relied upon by policymakers, regulators, institutions and the courts as AI adoption expands and matures.

This approach allows Mauritius to remain adaptive and forward-looking, avoiding premature or overly rigid regulation, while ensuring that appropriate safeguards and accountability mechanisms can be introduced as experience, evidence and capacity grow.

Relationship with Existing Legal and Institutional Guidelines

The FAIR Guidelines are intended to operate in harmony with existing legal and institutional arrangements, including data protection, consumer protection, administrative law and sectoral regulation.

Nothing in these Guidelines should be read as limiting or undermining existing legal obligations. Where overlaps arise, institutions and regulators are expected to apply the Guidelines in a manner that reinforces legal compliance, institutional integrity and public trust.

Scope Exclusion

These Guidelines do not apply to the development or use of Artificial Intelligence for military, defence, national security or counter-terrorism purposes. Such uses are governed by separate legal, constitutional and security frameworks and fall outside the scope of this document.

Guiding Values- The FAIR Pillars

Application of the FAIR Pillars

The FAIR pillars apply collectively and in balance. No single pillar should be considered in isolation and trade-offs may arise in practice. Where tensions occur, decisions should be guided by proportionality, public interest and the specific context in which an AI system is developed and used.

These pillars provide the interpretive lens through which the remaining sections of the Guidelines should be read and applied. They are intended to inform governance design, risk assessment, oversight mechanisms and future policy development.

Roles and Responsibilities - A Shared Governance Model

A Shared Approach to AI Governance

Effective AI governance requires shared responsibility. No single institution, sector or actor can manage the opportunities and risks of AI alone.

The FAIR Guidelines adopt a shared governance model, recognising the distinct but complementary roles of public institutions, regulators, the private sector, academia and civil society. This approach reflects the practical realities of AI development and use and supports collaboration while maintaining clear accountability.

Roles and responsibilities are expressed at a general level in these Guidelines. More detailed mandates, procedures and enforcement mechanisms will be defined through sectoral policies, regulatory instruments and legislation, as appropriate.

Government and Public Institutions

Government and public institutions play a central role in setting direction, establishing governance expectations and ensuring that AI adoption serves the public interest.

Key responsibilities include:

- Providing strategic leadership and policy coherence for AI adoption
- Ensuring that AI use in public administration is lawful, transparent and accountable
- Embedding the FAIR principles into public sector decision-making, procurement and service delivery

Building institutional capability and awareness across the public service

Supporting coordination across ministries, agencies and state-owned enterprises

Public institutions are expected to lead by example, particularly where AI systems affect rights, access to services or public trust.

Regulators and Oversight Bodies

Regulators and oversight bodies have a critical role in ensuring that AI systems operate within existing legal and regulatory guidelines.

Their responsibilities include:

- Interpreting and applying the FAIR Guidelines within sector-specific contexts
- Identifying high-risk or sensitive AI use cases requiring enhanced oversight
- Integrating AI considerations into existing supervisory, inspection and enforcement functions
- Coordinating with other regulators to avoid fragmentation and regulatory gaps

The FAIR Guidelines do not alter existing statutory mandates. Instead, it provides a common governance reference to support consistent and proportionate regulatory approaches.

Private Sector

Private sector organisations play a key role as developers, deployers, vendors and users of AI systems.

Their responsibilities include:

- Designing and deploying AI systems in alignment with the FAIR principles
- Understanding and managing risks associated with AI development and use
- Ensuring appropriate internal governance, documentation and oversight
- Cooperating with regulators and public institutions where required
- Acting responsibly when providing AI systems or services to others, including the government

Private sector responsibility applies regardless of organisational size. Expectations should be applied proportionately, recognising differences in capacity, scale and impact.

Academia and Research Institutions

Academic and research institutions contribute to the responsible development and use of AI through research, education and independent expertise.

Their roles include:

- Advancing knowledge on AI capabilities, limitations and impacts
- Supporting evidence-based policy and governance approaches
- Contributing to skills development and capacity-building
- Engaging in ethical and responsible research practices

Academic independence and integrity are essential to informed public debate and sound policy development.

Civil Society and Individuals

Civil society organisations, professional bodies and individuals play an important role in shaping public understanding and trust in AI.

Their contributions include:

- Providing feedback on the societal impacts of AI systems
- Supporting public awareness and digital literacy
- Representing the interests of affected communities
- Contributing to dialogue on ethical and responsible AI development and use

Meaningful engagement helps ensure that AI governance remains grounded in lived experience and public values.

Shared Responsibility and Collaboration

AI governance is most effective when responsibilities are shared, understood and coordinated.

All actors are encouraged to:

- Engage in constructive dialogue
- Share lessons learned and good practices
- Participate in collaborative initiatives where appropriate
- Contribute to continuous improvement of governance approaches

This shared governance model supports innovation while ensuring that accountability and public trust are maintained.

AI System Lifecycle - From Design to Decommissioning

Purpose of a Lifecycle-Based Approach

Artificial Intelligence systems are not static products. They evolve as data, models, use cases and operating contexts change. Risks and impacts may emerge gradually, rather than at the point of deployment.

For this reason, the FAIR Guidelines adopt a lifecycle-based approach to AI governance. This approach recognises that responsible AI development and use requires attention not only at the moment a system is introduced, but throughout its entire lifecycle.

Governance expectations should therefore be understood as ongoing obligations, rather than one-off checks or approvals.

Overview of the AI Lifecycle

For these Guidelines, the AI lifecycle is understood to include the following stages:



Not all AI systems will follow these stages linearly or formally. However, the lifecycle provides a common reference structure to support proportionate governance across sectors and use cases.

Conception and Design

The conception and design stage includes decisions about:

- the purpose of the AI system,
- the problem it seeks to address,
- the context in which it will be used,
- and the potential impact on individuals, organisations or society.

At this stage, organisations are expected to:

- consider whether AI is an appropriate solution,
- identify foreseeable risks and limitations,
- clarify accountability and ownership,
- and align system objectives with the FAIR principles.

Early consideration of purpose and impact helps prevent avoidable harm and reduces the need for corrective action later.

Data Collection and Preparation

Data is a critical component of most AI systems and a common source of risk.

During this stage, organisations should pay particular attention to:

- the lawful and ethical basis for data use,
- data quality, relevance and representativeness,
- potential sources of bias or exclusion,
- data protection and security requirements.

Where sensitive or personal data is involved, expectations regarding governance, oversight and safeguards increase accordingly.

Development and Testing

The development and testing stage includes model training, validation and performance assessment.

Good practice at this stage includes:

- testing systems under realistic conditions,
- assessing performance across relevant groups or scenarios,
- documenting assumptions, limitations and known risks,
- and ensuring that appropriate human oversight mechanisms are in place.

Testing should be proportionate to the system's intended use and potential impact. Higher-risk systems require more rigorous validation and documentation.

Deployment and Use

Deployment marks the point at which an AI system begins to influence real-world decisions or outcomes.

At this stage, organisations should ensure that:

- users understand the role of the AI system,
- responsibilities for oversight and intervention are clear,
- appropriate safeguards are active,
- and affected individuals are treated fairly and transparently.

Deployment should not be viewed as the end of governance responsibility, but as the start of active operational oversight.

Monitoring and Adaptation

Once in use, AI systems may change in performance or impact due to:
new data,

- changes in operating context,
- shifts in user behaviour,
- or evolving societal expectations.

Ongoing monitoring is therefore essential. Organisations should:

- track system performance and outcomes,
- identify unintended effects or emerging risks,
- review whether the system continues to meet its intended purpose,
- and adapt or intervene where necessary.

The intensity of monitoring should be proportionate to the system's impact and risk profile.

Decommissioning or Withdrawal

AI systems should not be assumed to operate indefinitely.

Decommissioning or withdrawal may be appropriate where:

- a system no longer serves its intended purpose,
- risks can no longer be effectively mitigated,
- legal or policy requirements change,
- or better alternatives become available.

Responsible decommissioning includes:

- managing data appropriately,
- documenting decisions and outcomes,
- and ensuring continuity or transition where services are affected.

Lifecycle Governance and Proportionality

Across all lifecycle stages, governance expectations should be risk-based and proportionate.

Low-impact AI systems may require minimal formal controls, while high-impact or sensitive systems warrant enhanced scrutiny, documentation and oversight. The lifecycle approach provides a structured way to apply FAIR principles consistently, without imposing unnecessary burdens.

Risk-Based Governance Model

Rationale for a Risk-Based Approach

Not all AI systems pose the same level of risk. The potential impact of an AI system depends on its purpose, context of use, degree of autonomy, scale of deployment and the nature of the decisions it influences.

The FAIR Guidelines, therefore, adopt a risk-based approach to AI governance. This approach ensures that governance expectations are proportionate, focusing attention and resources where risks are greatest, while avoiding unnecessary burdens for low-impact applications and startups.

A risk-based approach also supports innovation by allowing experimentation and learning, provided that appropriate safeguards are in place.

Understanding AI Risk in Context

Within these Guidelines, AI risk is understood as the likelihood and severity of harm that may arise from the design, deployment or use of an AI system.

Risk may relate to, among other things:

- impacts on individual rights and freedoms,
- fairness and discrimination,
- access to essential services,
- economic or financial harm,
- public safety,
- institutional trust and credibility.

Risk is context-dependent. The same AI system may present different risk profiles depending on how and where it is used. Governance decisions should therefore be informed by the use context rather than technology alone.

Key Risk Factors

When assessing risk, organisations and regulators should consider a combination of factors, including:

Impact on individuals or groups

Whether the AI system influences rights, opportunities, access to services or significant life outcomes.

Level of autonomy

The extent to which decisions are automated or can be overridden by human judgment.

Scale and reach

The number of people affected and the frequency of use.

Reversibility of outcomes

Whether errors or adverse outcomes can be easily corrected.

Sensitivity of data involved

Including personal, financial, health or otherwise sensitive data.

No single factor is determinative. Risk assessment should be based on the overall effect, informed by professional judgement.

Indicative Risk Tiers

For guidance purposes, AI systems may be understood as falling into broad, indicative risk categories, recognising that context matters and that boundaries are not always rigid. These categories are intended to help people think about potential impact and do not create legal classifications or prohibitions.

Low-risk systems

Low-risk AI systems are applications with limited impact, typically used to support routine tasks or everyday activities where errors are unlikely to cause harm or serious consequences.

These systems generally assist users by improving efficiency, organisation or convenience, without directly influencing significant decisions or outcomes. Where errors occur, their effects are usually minor and can be easily identified and corrected.

Illustrative examples of low-risk AI systems may include, but are not limited to:

- AI tools that assist with organising emails, documents or files
- AI systems that provide spelling, grammar or writing suggestions
- AI tools used for text translation or summarisation
- AI systems that support scheduling, planning or task management
- AI features that offer general recommendations or productivity support

In these contexts, basic safeguards such as user awareness, clear communication about system limitations and appropriate data handling practices are generally sufficient. Ongoing monitoring may still be appropriate, particularly where systems are used at scale or changes over time.

These examples are indicative only and do not constitute a legal classification. The level of oversight and governance applied should remain proportionate to the system's purpose, context and potential impact.

Medium-risk systems

Medium-risk AI systems are applications that influence decisions or processes in ways that may have noticeable effects on individuals, organisations or services, but where human judgement remains involved and outcomes can usually be reviewed, adjusted or corrected.

These systems often support decision-making rather than replace it. While they may shape priorities, recommendations or assessments, final decisions typically remain with human users. As a result, risks are more contained than in high-risk contexts, but still require appropriate care in design, deployment and use.

Illustrative examples of medium-risk AI systems may include, but are not limited to:

- AI tools used to support the shortlisting or screening of job applications
- AI systems that assist customer service or case-handling decisions
- AI tools that help prioritise cases, inspections or applications for review
- AI systems that support credit, loan or risk assessments, where final decisions are made by humans
- AI tools used to assist with scheduling, allocation of appointments or resource planning

In these contexts, safeguards such as clear user guidance, transparency about how AI outputs are generated, appropriate human oversight and the ability to review or override outcomes are particularly important. Monitoring system performance over time also helps ensure that impacts remain proportionate and aligned with organisational objectives.

These examples are indicative only and do not constitute a legal classification. The appropriate level of oversight and control should be determined based on context, scale and potential impact, taking into account how the AI system is used in practice.

High-risk systems

High-risk AI systems are applications used in sensitive or important contexts where errors, bias or misuse could lead to significant harm, unfair treatment, loss of rights, exclusion from essential services or serious loss of public trust.

These systems are often associated with decisions or recommendations that have a direct and meaningful impact on individuals or organisations. While such uses may offer benefits, they also require heightened care at the design, development and deployment stages, as mistakes or unintended effects can be difficult to reverse.

Illustrative examples of high-risk AI systems may include, but are not limited to:

- AI systems that influence eligibility for public services, social benefits or housing
- AI systems that affect access to education, employment or professional opportunities
- AI systems used in credit assessment, insurance or other financial decision-making
- AI systems that support or inform medical diagnosis, treatment or care decisions
- AI systems used in regulatory, compliance, inspection or enforcement activities
- AI systems that play a role in decisions affecting individual rights, obligations or legal status

For such systems, stronger safeguards are typically appropriate. These may include more robust testing and validation, clearer documentation, enhanced human oversight, mechanisms for review and redress and ongoing monitoring throughout the system's lifecycle.

These examples are indicative only and do not constitute a legal classification. The assessment of whether an AI system should be treated as high-risk depends on its purpose, context of use, scale and potential impact and should be informed by professional judgement and relevant policy or regulatory frameworks.

Unacceptable risk

Some uses of AI may raise fundamental concerns because of their potential impact on human dignity, fundamental rights or the public interest. These are situations where the use of AI may be considered inappropriate regardless of potential benefits.

The FAIR Guidelines do not define or prohibit such uses. Decisions about whether particular AI applications are unacceptable, restricted or prohibited are matters for law, policy and competent authorities and fall outside the scope of these Guidelines. This category is included to recognise that not all uses of AI are suitable in all contexts and that such judgements must be made through proper democratic and legal processes.

Illustrative examples of uses that may raise such concerns include, but are not limited to:

- the use of AI systems to enable indiscriminate or continuous mass surveillance in public or private spaces
- the use of AI to manipulate behaviour or choices in a manner that undermines individual autonomy or informed consent
- the use of AI to make irreversible or highly consequential decisions about individuals without meaningful human involvement or avenues for review
- the use of AI systems that intentionally discriminate against individuals or groups on prohibited grounds
- the use of AI in contexts where errors or misuse could result in serious harm to personal safety, dignity or fundamental rights, without adequate safeguards

These examples are indicative only and do not constitute a determination of legality or acceptability. Any assessment of such uses must be undertaken through appropriate legal, regulatory and policy processes.

Proportionate Governance Expectations

As risk increases, so too do expectations regarding governance, oversight and safeguards. In general terms:

- Low-risk systems may require minimal formal controls beyond basic transparency and good practice.
- Medium-risk systems warrant clearer documentation, oversight arrangements and performance monitoring.
- High-risk systems require enhanced scrutiny, stronger accountability, robust safeguards and ongoing review.

This graduated approach helps ensure that governance measures are effective, credible and manageable.

Role of Regulators and Institutions

Regulators and oversight bodies play an important role in interpreting and applying the risk-based approach within their respective sectors, taking into account the specific context, mandates and realities they operate in.

In practice, regulators may, as appropriate:

- identify AI use cases that warrant closer attention due to their potential impact,
- provide sector-specific guidance informed by these Guidelines,
- coordinate with other regulators where AI-related risks cut across institutional or sectoral boundaries and
- adapt their approach over time as experience, evidence and technology evolve.

In carrying out these roles, regulators may also take note of relevant regional and international developments related to AI governance and risk management. This includes ongoing work and discussions within regional and continental forums such as the Southern African Development Community (SADC), the African Union (AU) and the Common Market for Eastern and Southern Africa (COMESA), as well as other international initiatives.

Such engagement supports information-sharing, learning and coherence, while recognising that regulatory decisions in Mauritius remain grounded in national law, institutional mandates and domestic priorities.

The FAIR Guidelines support regulators by offering a shared conceptual reference for thinking about AI risks, while allowing sufficient flexibility to reflect sector-specific realities and evolving regional and international practice.

Continuous Risk Review

AI risk is not static. Systems may evolve over time and new risks may emerge as technologies and use patterns change.

Organisations are therefore expected to:

- periodically review risk assessments,
- update governance measures where necessary,
- and respond to emerging concerns in a timely manner.

This emphasis on continuous review reinforces the adaptive and responsible use of AI.

Escalation, Controls and Safeguards

Purpose of Escalation and Safeguards

A risk-based approach to AI governance is effective only if it is supported by clear escalation mechanisms and appropriate safeguards.

The purpose of escalation is not to slow innovation, but to ensure that as the potential impact of an AI system increases, so does the level of attention, oversight and control applied to it. Safeguards are intended to prevent harm, detect issues early and enable timely intervention where needed.

This section sets out the governance logic for escalation and safeguards. Detailed requirements will be defined through sectoral policies, guidance and legislation as appropriate.

Escalation Triggers

Escalation should occur when one or more factors indicate that an AI system may pose increased risk or sensitivity.

Indicative escalation triggers include:

- use of AI in decisions affecting rights, eligibility or access to essential services,
- increased system autonomy with limited human oversight,
- large-scale deployment affecting significant numbers of people,
- use of sensitive or high-risk data,
- material changes to system purpose, scope or performance,
- credible evidence of harm, bias or misuse.

Escalation does not imply prohibition. It signals the need for enhanced governance attention.

Governance Controls by Risk Level

Governance controls should scale with the assessed level of risk to be defined by a competent authority.

In broad terms, this may include:

For lower-risk systems

Basic documentation, transparency regarding use and routine operational oversight.

For medium-risk systems

Clear accountability arrangements, documented risk assessments, testing prior to deployment and defined oversight procedures.

For higher-risk systems

Enhanced scrutiny, formal approvals, robust documentation, stronger human oversight, ongoing monitoring and clearly defined escalation and intervention mechanisms.

Controls should be proportionate, practical and aligned with institutional capacity.

Human Oversight and Intervention

Human oversight remains an important safeguard in the development and use of Artificial Intelligence, particularly where AI systems influence decisions with meaningful consequences.

Depending on the context and level of risk, human oversight may involve:

- reviewing or validating AI-supported recommendations,
- overriding, suspending or correcting automated outputs and
- retaining human accountability for final decisions where AI is used to support or inform them.

The level and form of human oversight should be proportionate to Low-risk AI systems, used for routine or administrative purposes, may require limited or no active intervention beyond general supervision. Medium- and high-risk AI systems, by contrast, typically warrant clearer oversight arrangements and the ability for humans to intervene meaningfully when needed.

Human oversight and intervention do not make all uses of AI appropriate. Where the use of AI raises fundamental concerns for human dignity, rights or the public interest, such uses may be considered unacceptable regardless of the level of oversight applied.

Where human intervention or override occurs in medium- or high-risk contexts, these actions should be capable of being recorded and reviewed. This supports transparency, learning and accountability and helps organisations understand how AI systems are used in practice.

Oversight arrangements should be meaningful rather than symbolic. Where AI systems are used in high-impact contexts, organisations should ensure that those responsible for oversight have the authority, competence and access to information necessary to intervene effectively.

Human Oversight Models

Human oversight ensures that AI systems operate within defined boundaries and remain subject to human judgement and control.

Depending on context, oversight may take different forms, including:

- human review of AI-generated recommendations,
- human approval before automated actions are executed or
- ongoing human supervision of system behaviour.

Oversight arrangements should be clearly defined, practical and supported by appropriate authority and competence. Oversight that exists only in form, without real capacity to intervene, does not meet the intent of these Guidelines.

Transparency and Documentation

Transparency and documentation are essential safeguards, particularly as risk increases.

Organisations should be able to demonstrate:

- the purpose of an AI system,
- how it is intended to be used,
- key assumptions and limitations,
- known risks and mitigation measures,
- and who is responsible for oversight.

The depth of documentation should reflect the level of risk and impact. Excessive formality should be avoided for low-risk systems, while high-risk systems require robust and auditable records.

Monitoring, Incident Response and Corrective Actions

Safeguards must extend beyond deployment.

Organisations are expected to:

- monitor AI systems for performance and unintended effects,
- establish processes for identifying and responding to incidents,
- take corrective action where risks materialise,
- and, where necessary, suspend or withdraw systems.

Timely response to issues is essential to maintaining trust and preventing escalation of harm.

Prohibited and Restricted Uses

Certain uses of AI may be inappropriate or unacceptable in specific contexts, particularly where they conflict with fundamental rights, legal obligations or public interest considerations.

The FAIR Guidelines do not establish prohibitions. However, they recognise that:

- some uses may warrant restriction or prohibition through law or regulations;
- such determinations should be evidence-based and proportionate; and
- they should be made transparently through appropriate democratic and regulatory processes.

These Guidelines provide the conceptual basis for such decisions, should they be required in the future.

Balancing Innovation and Safeguards

Escalation and safeguards should support responsible innovation, not discourage it.

Where risks are well understood and managed, organisations should be encouraged to experiment, learn and improve. Where uncertainty or potential harm is high, caution and stronger controls are justified.

This balance is central to the FAIR approach.

Data Governance and Data Protection Alignment

Importance of Data Governance in AI

Data is a foundational component of most AI systems and a primary source of both value and risk. The quality, legality and governance of data directly influence the reliability, fairness and trustworthiness of AI outcomes.

Effective AI governance therefore depends on sound data governance practices that are consistent with existing legal obligations, institutional responsibilities and public expectations. Poor data practices can undermine even well-designed AI systems.

This section sets out baseline expectations to ensure that AI systems are supported by data practices that are lawful, appropriate and aligned with public interest considerations.

The FAIR Guidelines operates in alignment with existing data protection, privacy and confidentiality obligations applicable in Mauritius.

Where AI systems involve personal or sensitive data, organisations are expected to:

- comply fully with applicable data protection legislation,
- respect principles of lawfulness, fairness and purpose limitation,
- ensure transparency regarding data use,
- and apply appropriate safeguards to protect individuals' rights.

Nothing in this Guidelines reduces or replaces existing legal obligations. Rather, it reinforces the expectation that AI systems should be designed and used in a manner that strengthens compliance and institutional trust.

Data Quality, Relevance and Representativeness

AI systems rely on data that is fit for purpose.

Organisations should take reasonable steps to ensure that:

- data used is relevant to the intended purpose,
- data quality is sufficient to support reliable outcomes,
- known limitations or gaps are understood and documented,
- and potential sources of bias or distortion are identified and addressed.

Where data reflects historical patterns or systemic inequalities, organisations should consider how this may affect AI outputs and whether mitigation measures are appropriate, particularly in higher-impact contexts.

Lawful and Responsible Data Use

AI development and deployment should be grounded in clear and lawful bases for data use.

This includes:

- respecting data minimisation principles,
- avoiding secondary uses incompatible with original purposes,
- ensuring appropriate consent or legal authority where required,
- and applying safeguards for sensitive categories of data.

Responsible data use also requires judgement. Even where data use is lawful, organisations should consider whether it is appropriate and proportionate in light of the AI system's purpose and potential impact.

Cross-Border Data Considerations

Mauritius is a small, open economy with strong international linkages. AI systems may involve data storage, processing or model development across borders, including through cloud services, external vendors or regional partnerships.

In such cases, organisations should:

- be mindful of applicable cross-border data transfer and data protection obligations,
- ensure that appropriate safeguards and protections travel with the data and
- consider risks related to dependency, oversight, accountability and jurisdiction.

These considerations are increasingly relevant within the African context. Across the continent, initiatives led by the African Union and regional bodies are promoting stronger cooperation on data protection, digital governance and trusted cross-border data flows, as part of broader efforts to support digital transformation and economic integration.

Strategic decisions regarding cross-border data use should therefore balance openness to international and regional collaboration with the need to protect public interest, maintain accountability and support long-term national resilience.

Data Governance Across the AI Lifecycle

Data governance responsibilities extend across the entire AI lifecycle.

Organisations should consider:

- data governance at the design stage, including data sourcing decisions,
- safeguards during development and testing,
- controls during deployment and operation,
- and appropriate data handling during system modification or decommissioning.

Lifecycle-based data governance helps ensure that risks are managed proactively rather than reactively.

Shared Responsibility for Data Governance

Data governance in AI systems often involves multiple actors, including data providers, system developers, deployers and users.

Responsibilities should be clearly understood and documented, particularly where:

- data is shared between organisations,
- third-party datasets or models are used,
- or AI systems are procured from external vendors.

Clear allocation of responsibility supports accountability and reduces governance gaps.

Proportionality and Practicality

Data governance expectations should be proportionate to risk and context.

Low-risk AI systems may require limited formal controls, while high-impact systems warrant stronger governance, documentation and oversight. The objective is not to impose unnecessary burdens, but to ensure that data practices support responsible and trustworthy AI use.

Transparency, Explainability and Human Oversight

Purpose and Rationale

Transparency, explainability and human oversight are essential to maintaining trust in the use of Artificial Intelligence, particularly where AI systems influence decisions that affect individuals, organisations or public interests.

These elements help ensure that AI systems are not treated as opaque or unquestionable tools and that their use remains understandable, reviewable and accountable within existing institutional and legal guidelines.

Transparency of AI Use

Organisations using AI systems should be transparent about when and how AI is used, particularly in contexts where outcomes may affect rights, access to services or significant interests.

Transparency may include:

- informing users or affected parties that AI is being used,
- describing the role played by the AI system in decision-making,
- clarifying whether decisions are fully automated or supported by AI.

The level of transparency should be proportionate to context and risk. Transparency expectations increase where AI use has material consequences or where individuals may reasonably expect explanation.

Explainability and Understandability

Explainability refers to the ability to provide meaningful information about how an AI system is developed, how it operates and how it contributes to outcomes in practice. Understandability means that explanations are clear and meaningful to the people who need to rely on them, whether they are system users, decision-makers, regulators or affected individuals. The objective is not to require disclosure of proprietary algorithms, source code or models. Rather, it is to ensure that:

- decision logic can be explained at an appropriate level,
- key factors influencing outcomes can be identified and
- relevant limitations, uncertainties or assumptions are acknowledged.

Expectations around explainability and understandability should be proportionate and reflect:

- the nature of the AI system,
- its intended use and
- the potential impact of its outcomes.

In higher-impact contexts, explanations should be sufficient to support review, challenge and correction where appropriate, without undermining legitimate intellectual property or commercial interests.

Authority and Competence

Human oversight is effective only where those responsible:

- understand the purpose, functioning and limitations of the AI system,
- have access to information necessary to perform oversight and
- possess the authority to intervene, escalate concerns or suspend use where required.

Organisations should ensure that oversight responsibilities are assigned to individuals or roles with appropriate skills, training and institutional backing. Where assurance, audit or review of AI systems is required, this may involve access to relevant documentation, testing results or system behaviour under controlled and confidential conditions, without requiring disclosure of proprietary source code.

Transparency Across the AI Lifecycle

Transparency and explainability considerations apply throughout the AI lifecycle.

This includes:

documenting design choices and assumptions,
recording changes to system behaviour or scope
and maintaining records that support later review or audit.

Clear records across the AI lifecycle support responsibility, learning and continuity, even as staff or systems change.

Balancing Transparency with Practical Constraints

Transparency and explainability should be pursued in a manner that is practical and proportionate.

In some cases, full transparency may be limited by:

- technical complexity,
- intellectual property considerations,
- or security concerns.

Where limitations exist, organisations should aim to provide the best possible level of explanation consistent with context, risk and legitimate constraints.

Contribution to Trust and Accountability

By promoting transparency, explainability and meaningful human oversight, the FAIR Guidelines reinforce public trust, support institutional accountability and enables responsible AI use across sectors.

These elements help ensure that AI systems support human and institutional decision-making, rather than replacing it

Accountability, Auditability and Record-Keeping

Purpose and Importance

Accountability is a central pillar of the FAIR Guidelines. As AI systems increasingly influence decisions, processes and outcomes, responsibility must remain clearly assigned, visible and enforceable.

Auditability and record-keeping support accountability by ensuring that decisions related to AI systems can be understood, reviewed and, where necessary, challenged. Without appropriate records, meaningful oversight is not possible.

Clear Assignment of Responsibility

For every AI system in use, organisations should ensure that responsibility is explicitly assigned.

This includes clarity regarding:

- who is responsible for authorising the use of the AI system,
- who oversees its operation and performance,
- and who is accountable for addressing risks, errors or harm.

Responsibility should not be diffused across systems, vendors or automated processes. Where multiple parties are involved, roles and responsibilities should be clearly documented and understood.

Accountability Across the AI Lifecycle

Accountability applies throughout the AI lifecycle, not only at the point of deployment.

Organisations are expected to maintain accountability during:

- design and development,
- testing and validation,
- deployment and use,
- ongoing monitoring and modification,
- and decommissioning or withdrawal.

Changes to system purpose, scope or behaviour should trigger a review of accountability arrangements to ensure they remain appropriate.

Auditability and Traceability

Auditability refers to the ability to examine how an AI system was designed, how it operates and how specific outcomes were produced.

To support auditability, organisations should be able to:

- trace key decisions and changes made throughout the lifecycle,
- understand data sources and model versions used,
- reconstruct decision pathways where required,
- and demonstrate compliance with governance expectations.

The depth of auditability required should be proportionate to the system's risk and impact.

Record-Keeping Practices

Appropriate record-keeping is a practical requirement for accountability.

Records may include, as relevant:

- system purpose and intended use,
- risk assessments and mitigation measures,
- testing and validation results,
- oversight arrangements and escalation processes,
- incidents, complaints and corrective actions,
- significant updates or modifications.

Record-keeping should be fit for purpose, avoiding unnecessary formality for low-risk systems while ensuring robustness for higher-risk applications.

Use of Records for Oversight and Learning

Records should not be treated as a compliance burden alone. They serve important functions in:

- supporting internal governance and decision-making,
- enabling regulatory or audit review where required,
- facilitating learning and improvement over time,
- and strengthening institutional memory.

Well-maintained records help organisations identify patterns, improve practices and respond more effectively to emerging risks.

Accountability in Multi-Party and Vendor Contexts

Where AI systems involve multiple organisations, including external vendors or service providers, accountability arrangements should be carefully managed.

Organisations procuring or using AI systems should:

- understand their own responsibilities,
- avoid assumptions that accountability rests solely with vendors,
- ensure that contractual arrangements support transparency and oversight,
- and retain sufficient control to meet governance expectations.

Shared responsibility does not mean unclear responsibility.

Proportional and Practical Application

Accountability, auditability and record-keeping should be applied proportionately and practically. The objective is to ensure meaningful oversight and responsibility, not to impose excessive administrative burden. Expectations should reflect context, capacity and risk, while remaining credible and enforceable.

Procurement, Third-Party and Vendor Governance

Importance of Procurement in AI Governance

A significant proportion of AI systems used in Mauritius is likely to be procured from external vendors, integrated into existing platforms or delivered as part of outsourced services.

Procurement decisions therefore play a critical role in shaping how AI is designed, deployed and governed. Weak procurement practices can undermine accountability, transparency and public trust, regardless of how well governance principles are articulated elsewhere.

This section sets out baseline governance expectations for procuring and managing AI systems supplied by third parties.

Responsibility Does Not Transfer with Procurement

Procuring an AI system from a third party does not transfer responsibility for its use or impact.

Organisations deploying AI systems remain responsible for:

- ensuring alignment with the FAIR principles,
- understanding how the system is intended to function,
- overseeing its use in context,
- and addressing risks or harm that may arise.

Vendor involvement does not diminish the obligation to exercise due care and judgement.

Due Diligence and Risk Awareness

Before procuring or deploying an AI system, organisations should undertake appropriate due diligence, proportionate to the system's risk and intended use.

This may include consideration of:

- the purpose and limitations of the system,
- data sources and data handling practices,
- governance and oversight mechanisms,
- known risks, biases or constraints,
- and the vendor's capacity to support transparency and accountability.

For higher-risk systems, more structured due diligence is warranted.

Transparency and Information from Vendors

Effective governance requires access to sufficient information about procured AI systems.

Organisations should seek, where appropriate:

- clear descriptions of system functionality and intended use,
- information on data inputs and model behaviour at a meaningful level,
- documentation to support oversight, audit or review,
- clarity on update mechanisms and system changes.

The objective is not full technical disclosure, but practical transparency that enables responsible use.

Contractual and Governance Considerations

Contracts involving AI systems should support governance objectives.

Where appropriate, contractual arrangements may address:

- roles and responsibilities for oversight and support,
- obligations relating to transparency and cooperation,
- processes for incident reporting and remediation,
- expectations regarding system updates or modifications,
- and termination or withdrawal where risks cannot be managed.

Contracts should not be used to shift responsibility away from deploying organisations.

Ongoing Vendor Management

Governance responsibilities continue after procurement.

Organisations should:

- monitor system performance and behaviour over time,
- remain alert to changes introduced through updates or new features,
- review whether the system continues to meet governance expectations,
- and engage with vendors where issues arise.

Vendor management should be proportionate and risk-informed, recognising capacity constraints while maintaining accountability.

Public Sector Considerations

In the public sector, procurement of AI systems carries additional responsibilities due to the use of public resources and the impact on citizens.

Public institutions should ensure that AI procurement:

- aligns with public service values,
- supports transparency and accountability,
- and reinforces public trust.

Where AI systems are used in sensitive or high-impact public functions, enhanced scrutiny and oversight are justified.

Supporting Innovation through Responsible Procurement

Responsible procurement does not mean risk avoidance.

Well-designed procurement processes can:

- encourage innovation,
- support local and emerging providers,
- and promote responsible AI practices.

The FAIR Guidelines support procurement approaches that balance innovation with governance, ensuring that AI adoption delivers value while managing risk.

Generative, Foundation and Agentic AI Systems

Context and Rationale

Recent advances in Artificial Intelligence have led to the widespread availability of generative, foundation and agentic AI systems, including systems capable of producing text, images, code, recommendations and actions with varying degrees of autonomy.

These systems present significant opportunities for productivity, creativity and service innovation. At the same time, they introduce distinct governance considerations that differ in important ways from more traditional, task-specific AI applications.

The FAIR Guidelines recognise the need to address these developments explicitly, while avoiding premature or overly rigid responses.

Distinctive Characteristics

Generative, foundation and agentic AI systems often share characteristics that warrant particular attention, including:

- broad and general-purpose capabilities,
- use across multiple contexts and sectors,
- dependence on large and diverse datasets,
- probabilistic outputs rather than deterministic results,
- potential for content generation at scale,
- and, in some cases, increasing degrees of autonomy.

These characteristics can amplify both benefits and risks, particularly when such systems are embedded in decision-making processes or public-facing services.

Governance Considerations

When developing, deploying or using these systems, organisations should give due consideration to issues such as:

- accuracy and reliability of outputs,
- risks of misleading, fabricated or harmful content,
- intellectual property and ownership concerns,
- transparency regarding AI-generated content,
- appropriate human oversight and intervention,
- and alignment with existing legal and ethical obligations.

These considerations apply regardless of whether such systems are developed locally or accessed through external platforms or services.

Context-Sensitive Application

The governance expectations for generative and agentic AI systems should be context sensitive.

Low-risk uses, such as internal productivity tools or creative experimentation, may require limited formal controls. Higher risk uses, particularly those affecting rights, access to services, public information or institutional decision-making, warrant heightened scrutiny and safeguards.

Organisations should avoid deploying such systems in high-impact contexts without adequate understanding of their limitations and risks.

Responsibility for Outputs and Use

The use of generative or agentic AI systems does not diminish responsibility for outcomes.

Organisations and individuals remain responsible for:

- how these systems are used,
- how outputs are interpreted or acted upon,
- and how risks are managed in practice.

Automated generation does not absolve users or institutions of accountability.

Continuous Learning and Adaptation

Given the pace of change in this area, governance approaches should be adaptive.

Institutions are encouraged to:

- monitor emerging risks and best practices,
- share lessons learned,
- and update internal policies and controls as experience grows.

This adaptive approach supports innovation while maintaining responsible use.

Sectoral Adaptation and Implementation Principles

Purpose of Sectoral Adaptation

Artificial Intelligence is used in diverse contexts across the economy and society. The risks, opportunities and governance needs associated with AI in financial services, healthcare, education, tourism, logistics or public administration are not identical.

The FAIR Guidelines are therefore designed to operate as a horizontal governance baseline, while allowing for sector-specific adaptation where justified by context, risk and regulatory mandate.

This approach ensures coherence across sectors without imposing a one-size-fits-all model.

Role of Sector Regulators and Lead Institutions

Sector regulators and lead institutions play a central role in translating the FAIR Guidelines into practical, context-aware governance arrangements.

Within their respective mandates, regulators are expected to:

- interpret the FAIR principles in light of sector-specific risks and practices,
- identify AI use cases that warrant enhanced oversight,
- issue guidance or requirements consistent with this Guidelines,
- and coordinate with other regulators where AI systems cut across sectors.

Nothing in these Guidelines alters existing statutory responsibilities. Rather, it provides a shared reference point to support consistency and proportionality.

Consistency and Avoidance of Fragmentation

While sectoral adaptation is necessary, excessive divergence can create confusion, compliance burdens and governance gaps.

Sector-specific measures should therefore:

- remain aligned with the FAIR pillars and risk-based logic,
- avoid conflicting or duplicative requirements,
- and be communicated clearly to affected organisations.

Coordination across sectors is particularly important where AI systems are deployed across multiple domains or markets.

Proportionality and Capacity Considerations

Sectoral implementation of FAIR should reflect institutional capacity and maturity.

Regulators and institutions should:

- apply governance expectations proportionately,
- recognise differences in organisational size and capability,
- and avoid imposing requirements that are impractical or disproportionate to risk.

This is especially relevant for small and medium-sized enterprises and emerging sectors, where excessive compliance burdens could stifle innovation.

Supporting Innovation within Sectors

Sectoral adaptation should not focus solely on risk mitigation.

Well-designed governance approaches can:

- support responsible experimentation,
- encourage adoption of beneficial AI applications,
- and create clarity for innovators and investors.

The FAIR Guidelines supports sectoral approaches that balance innovation and safeguards, enabling AI to contribute meaningfully to sectoral development goals.

Learning and Cross-Sector Exchange

AI governance is an evolving field. Sectors are encouraged to:

- share experiences and lessons learned,
- contribute to common understanding of emerging risks,
- and participate in cross-sector dialogue.

Such exchange supports continuous improvement and reduces the likelihood of fragmented or inconsistent governance practices.

Capacity Building and Competence

Why Capacity Matters

Effective AI governance depends not only on principles and guidelines, but on the capacity of institutions and people to understand, apply and oversee AI systems in practice.

Without adequate skills, awareness and organisational readiness, even well-designed governance arrangements risk remaining formal or symbolic. Capacity building is therefore a foundational enabler of the FAIR Guidelines.

Institutional Capacity in the Public Sector

Public institutions play a critical role in setting standards, procuring AI systems and overseeing their use.

Public sector capacity building should focus on:

- improving understanding of AI capabilities and limitations,
- strengthening the ability to assess risk and impact,
- supporting informed decision-making and oversight,
- and ensuring that accountability arrangements are meaningful.

This does not require every institution to become technically specialised. Rather, it requires sufficient institutional literacy to engage responsibly with AI systems and external providers.

Skills and Competence in Organisations

Organisations deploying AI systems should ensure that relevant staff possess the competence required to fulfil their roles under the FAIR Guidelines.

This may include:

- understanding the purpose and functioning of AI systems in use,
- recognising potential risks and limitations,
- knowing when and how to escalate concerns,
- and being able to exercise effective human oversight.

Competence requirements should be proportionate to responsibility and risk and may evolve over time as AI use matures.

Professional Responsibility and Ethics

Where AI systems are used in professional contexts - such as public administration, finance, healthcare, education or advisory services - professional responsibility remains paramount.

AI should support, not replace, professional judgement. Organisations and professionals are expected to:

- use AI systems responsibly,
- remain accountable for decisions and outcomes,
- and adhere to applicable professional standards and codes of conduct.

Supporting SMEs, Start-ups and Technopreneurs

Small and medium-sized enterprises and new adopters of AI may face capacity constraints.

Governance approaches should therefore:

- avoid unnecessary complexity,
- provide clear guidance and support,
- and encourage responsible experimentation.

Building capacity across the economy supports inclusive and sustainable AI adoption, consistent with the FAIR principles.

Continuous Learning and Adaptation

AI technologies, use cases and governance expectations will continue to evolve.

Institutions and organisations are encouraged to:

- invest in ongoing learning and skills development,
- stay informed of emerging risks and best practices,
- and adapt governance arrangements as experience grows.

Capacity building should be understood as an ongoing process, not a one-time exercise.

Monitoring, Review and Continuous Improvement

Purpose and Rationale

Artificial Intelligence technologies, use cases and risks evolve over time. Effective governance must therefore be capable of learning, adjustment and improvement.

The FAIR Guidelines are designed to be a living governance instrument, supported by regular monitoring and periodic review. This ensures that governance approaches remain relevant, proportionate and aligned with both national priorities and international developments.

Monitoring of AI Use and Governance Practices

Institutions and organisations using AI systems are encouraged to monitor:

- how AI systems perform in practice,
- whether systems continue to serve their intended purpose,
- the emergence of unintended effects or risks,
- and the effectiveness of existing safeguards and oversight mechanisms.

Monitoring should be proportionate to risk and impact. High-impact or sensitive AI systems warrant more structured and frequent monitoring, while low-risk applications may require only routine review.

Learning from Experience

Monitoring should support learning, not blame.

Organisations and institutions are encouraged to:

- reflect on what has worked well and what has not,
- identify recurring challenges or gaps,
- and share lessons learned where appropriate.

Learning from experience helps improve governance practices over time and supports more informed policy and regulatory development.

Periodic Review of the FAIR Guidelines

The FAIR Guidelines itself should be reviewed as and when required to ensure it remains fit for purpose.

Reviews may consider:

- technological developments and emerging AI capabilities,
- changes in patterns of AI use across sectors,
- evidence of new or evolving risks,
- international policy and regulatory developments,
- and feedback from stakeholders.

Periodic review allows the Guidelines to evolve without requiring frequent or disruptive changes to underlying legislation or institutional arrangements.

Adaptation and Continuous Improvement

Where monitoring or review indicates that governance measures are insufficient or misaligned, adjustments should be made in a timely and proportionate manner.

Adaptation may take the form of:

- updated guidance or playbooks,
- revised sectoral approaches,
- strengthened oversight mechanisms,
- or, where justified, new policy or regulatory measures.

Continuous improvement supports responsible AI adoption while preserving flexibility and innovation.

Maintaining Trust and Credibility

Regular monitoring, review and adaptation reinforce public trust in AI governance. By demonstrating that AI use is subject to ongoing scrutiny and improvement, institutions signal that governance is active, responsible and responsive, rather than static or symbolic.

From Guidelines to Policy and Legislation

Purpose of This Section

The FAIR Guidelines are designed to serve as a foundation for future policy, regulatory and legislative instruments, developed progressively and proportionately as Artificial Intelligence adoption matures in Mauritius.

This section clarifies how the Guidelines are intended to inform and support that evolution, while preserving flexibility and institutional discretion.

FAIR as a Policy Reference Point

The FAIR Guidelines provide a common governance reference for policymakers, regulators and institutions when developing AI-related policies and guidance.

Elements of the Guidelines may be translated into:

- government policies and administrative instructions,
- sectoral guidelines and codes of practice,
- procurement standards and contractual clauses,
- supervisory expectations and regulatory guidance.

Not all elements of FAIR are intended to become binding. The Guidelines distinguish between foundational principles and context-dependent governance mechanisms, allowing policymakers to determine what level of formalisation is appropriate in each case.

Respect for Institutional Roles and Legal Processes

The FAIR Guidelines do not pre-empt the role of Parliament, regulators or the courts.

Decisions regarding:

- the form of regulation,
- the allocation of statutory powers,
- enforcement mechanisms,
- and legal remedies


remain matters for established constitutional and legal processes.

The Guidelines are intended to support decision-making within the prescribed framework, not to substitute any democratic or legal authority.

Concluding Note

The FAIR Guidelines reflect the commitment of the Ministry to a measured, responsible and forward-looking approach to Artificial Intelligence.

They provide a strong foundation to guide AI action today and flexible enough to fuel the socio-economic growth aspirations through AI. In doing so, it supports the development and use of responsible and trustworthy AI.



FAIR

**GUIDELINES FOR
THE DEVELOPMENT
AND USE OF
ARTIFICIAL
INTELLIGENCE**