



AI Workplace and Job Transformation Guidelines

Date: June 2026



◆ Disclaimer/Legal Rights

The Ministry of Communications and Information Technology (MCIT), has designed and created this publication, entitled AI Workplace and Job Transformation Guidelines, reference DIA-GUI-004 (hereinafter referred to as the “Work”), primarily as a resource for public and private sector organizations in Qatar adopting AI.

The “Work” has been prepared in accordance with the laws of the State of Qatar and does not confer, and may not be used to support, any right on behalf of any person or entity against the State of Qatar or its agencies or officials. If a conflict arises between this document and the laws of Qatar, the latter shall take precedence. Every effort has been made to ensure the “Work” is accurate, but no warranty, guarantee or undertaking is given regarding the accuracy, completeness or currency of the “Work”. Links to other websites are inserted for convenience only and do not constitute endorsement of material at those sites or any associated organization, product or service.

Any reproduction of this “Work” either in part or in full and irrespective of the means of reproduction, shall acknowledge MCIT as the source and owner of the “Work”. Any reproduction concerning the “Work” with intent of commercialization shall seek a written authorization from MCIT. MCIT reserves the right to assess the functionality and applicability of all such reproductions developed for commercial intent. The authorization from MCIT shall not be construed as an endorsement of the developed reproduction and the developer shall in no way publicize or misinterpret this in any form of media or personal / social discussions.

◆ Legal Mandate

The guidelines have been issued pursuant to Amiri Decision No. (57) of 2021 regarding the specification of ministries’ mandates, which includes the competencies of the Ministry of Communications and Information Technology. It is also issued under Ministerial Decision No. (17) of 2024, which established departments within the administrative units of the Ministry of Communications and Information Technology and defined their functions.

Strategic Alignment



Qatar National Vision 2030

- Foster innovation in a **knowledge-based economy**
- Leverage **advanced technologies** to support **sustainable development** and **economic growth**



Third National Development Strategy (2024-2030)

- The strategy positions information technology and digital capacity as a **core enabling sector for economic diversification** and for building Qatar's long-term capabilities in areas like AI and other emerging technologies.
- The strategy seeks to establish Qatar as a **smart, technologically sophisticated nation**, and it recognizes that success hinges on a digitally prepared society.



Digital Agenda 2030

- Build a **digitally savvy population**; nurture talent; provide access to training
- Drive **digital adoption** for diversification and competitiveness



NCSA's Secure Usage and Adoption of AI

- Support the **secure and responsible adoption** of AI across organizations in Qatar.
- Encourage organizations to consider **cybersecurity, information security, access control, monitoring, incident response and safe-use controls** when introducing AI into workplace processes and job transformation initiatives.



Document Summary

Name	AI Workplace and Job Transformation Guidelines
Version	1.0
Document Reference	DIA-GUI-004
Document Type	Guidelines
Summary	These Guidelines provide a human-centered, task-based framework for redesigning jobs and managing workforce transformation in the age of artificial intelligence. They guide organizations in identifying which tasks can be AI-augmented, which should remain human-led, and how to recombine work into future-ready roles. Structured in two phases - (1) task-based job redesign and (2) skills empowerment and upskilling - the Guidelines translate Qatar's national digital ambitions into a practical method that organizations can apply across sectors. They emphasize feasibility ("Can AI help?"), appropriateness ("Should AI help?"), and clear human-in-the-loop controls to ensure quality, trust, and accountability. Intended for public and private sector leaders, HR and transformation teams, and operational managers, the document provides a disciplined approach to adopting AI that protects people, strengthens skills, and enhances productivity.
Applicable To	• Recommendation for all private sector organizations • Recommendation for all public-sector institutions, including government and semi-government entities, state-owned enterprises and specialized agencies
Publishing Date	June 2026
Owner	MCIT



◆ Acronyms

Abbreviation	Full Form
AI	Artificial Intelligence
API	Application Programming Interface
APQC	American Productivity & Quality Center (referenced via “APQC Process Classification Framework”)
DA2030	Digital Agenda 2030
HR	Human Resources
IoT	Internet of Things
ISO/IEC	International Organization for Standardization / International Electrotechnical Commission (mentioned as “ISO/IEC process modelling guidelines”)
IT	Information Technology
KPI	Key Performance Indicator
MCIT	Ministry of Communications and Information Technology
NCSA	National Cyber Security Agency
NPC	National Planning Council
OECD	Organisation for Economic Co-operation and Development
PCF	Process Classification Framework (from “APQC Process Classification Framework (PCF)”)
QCB	Qatar Central Bank
QNV2030	Qatar National Vision 2030
QRDI	Qatar Research, Development and Innovation Council
SOP	Standard Operating Procedure
STEM	Science, Technology, Engineering and Mathematics
TASMU	TASMU Smart Qatar Program
WEF	World Economic Forum



◆ Glossary Of Terms And Definitions

Term	Full Form
AI augmentation	Use of AI to support, accelerate, or improve human tasks while keeping people responsible for outcomes.
Task-based job redesign	Breaking a role into specific tasks to assess how work should change with AI.
Can score	A feasibility score showing whether AI can practically support a task based on data, rules, and system readiness.
Should score	An appropriateness score showing whether AI should support a task based on risk, judgment, trust, and learning value.
Human-in-the-loop	A control approach where a person reviews, approves, edits, or overrides AI outputs before final action.
Augment now	Tasks that are feasible and appropriate for immediate AI support.
Augment soon	Tasks that may benefit from AI once data, processes, integrations, or controls improve.
Keep human-led	Tasks that should remain primarily performed by people due to judgment, empathy, safety, accountability, or trust.
Agentic AI	AI systems that can perform multi-step actions across tools or systems with limited human intervention.
Skills heatmap	A visual matrix comparing future-role skill requirements against current employee or team proficiency levels.



Table of Contents

01	◆	Introduction	8
02	◆	Scope and Application	12
03	◆	Phases of AI Workforce and Job Transformation	13
		3.1 Phase 1 Redesign the Work	14
		3.2 Phase 2 Empower and Upskill	25
04	◆	Document Control	28

01 Introduction

Innovation is an indelible imprint of modern capitalist societies where firms, big and small, compete by adopting new technologies to lower costs and maximize revenues, making technical change a key driver of business success. Historically, industrial revolutions have reshaped society, from the steam engine in the 17th century to today's Fourth Industrial Revolution, driven by rapid advances in information and communication technologies. These include the **internet, artificial intelligence (AI), machine learning, data analytics, robotics, blockchain, and the Internet of Things (IoT)**, each of which is reshaping how economies produce, distribute, and consume value.

Technical change often involves replacing human labor with advanced technologies, leading to automation. In the context of AI, these risks can become more complex because AI systems may not only automate routine tasks but also support or influence judgment-based activities such as drafting, analysis, prioritization, decision support and customer interaction. If adopted without clear controls, AI can contribute to task displacement, workforce anxiety, skills polarization, over-reliance on automated outputs, unclear accountability, biased or inconsistent outcomes, and reduced transparency for employees and service users. These risks need to be managed across people, process and technology dimensions: people need the skills and confidence to work with AI responsibly; processes need clear decision rights, escalation points and quality controls; and technology needs appropriate security, monitoring and governance safeguards. While this boosts productivity, many workers may face technological unemployment (jobs that are no longer needed) or experience deskilling, where their existing skills become outdated and less valuable in the evolving labor market.

AI is already reshaping jobs and workplaces, with the potential for more significant changes than seen in past industrial revolutions. Tasks once considered resilient to technical change such as complex information processing, knowledge work and high-order cognitive activities are now squarely in scope.

In Qatar, the stakes are especially high because AI adoption will directly affect the country's ability to build a competitive, diversified and knowledge-based economy while protecting the quality, resilience and inclusiveness of its workforce. If AI is adopted without clear workforce planning, organizations may face unmanaged task displacement, widening skills gaps, over-reliance on automated outputs, inconsistent service quality and reduced trust among employees and service users. These risks are particularly important in Qatar because many services are delivered across interconnected public and private sector ecosystems, often in bilingual environments, where changes to roles, skills and accountability can have cross-sector impact.

The **National Development Strategy 3 (NDS-3)**, the **Digital Agenda 2030 (DA2030)** and the broader **Qatar National Vision 2030 (QNV 2030)** converge on a single imperative: build a competitive, diversified, knowledge-based economy powered by data and advanced analytics. The State of Qatar has invested heavily in fiber connectivity, cloud capacity and smart-city platforms through TASMU; yet the most valuable asset remains its people. The question, therefore, is no longer whether organizations should adopt AI, but how they can do so in a way that unlocks productivity and innovation while protecting and elevating the human contribution.

1.1 | What is the purpose of these Guidelines?

These Guidelines respond to the workplace transformation challenge created by AI. Their purpose is to **help public and private sector organizations in Qatar adopt AI in a way that improves productivity, service quality and innovation while protecting people, job quality, trust and accountability.**

The Guidelines provide a practical method for:

- Diagnosing organizational readiness.
- Redesigning work at the level of individual tasks.
- Identifying where AI can and should support work.
- Defining human-in-the-loop controls.
- Recombining tasks into future roles, and linking those roles to skills development and continuous learning.

The Guidelines translate Qatar's national digital ambitions into a practical roadmap that employers can apply, whether a ministry is digitizing citizen services or a private-sector firm is improving back-office operations. In addition, they embed principles of fairness, inclusion, transparency and human oversight so that AI becomes a catalyst for better jobs rather than a source of unmanaged displacement, inequality or loss of accountability.

While the document is written primarily for chief executives, public-sector leaders, HR directors and transformation managers, it should also serve professional bodies, training providers and regulators who share responsibility for a smooth and just transition. Above all, it is intended for the employees whose day-to-day experience of work will change most visibly, such as:

- Engineers will increasingly oversee AI-enabled digital twins and automated control systems rather than carrying out routine field checks.
- Healthcare professionals will review AI-assisted diagnostics and personalized treatment recommendations instead of relying only on manual test results.
- Financial analysts and auditors will validate AI-generated risk models and anomaly detections rather than compiling spreadsheets line by line.
- Service advisors will manage conversational AI agents and multilingual assistants, stepping in for complex or sensitive cases.
- Creative and professional roles will evolve, with staff refining AI-drafted reports, contracts, or design concepts instead of starting from a blank page



1.2 | How are these guidelines aligned with international best practices?

The AI Workplace and Job Transformation Guidelines aligns with international frameworks such as the OECD's AI Principles¹ and the World Economic Forum's human-centric reskilling agenda, which both emphasize that effective AI transformation starts with understanding work at the task level and then linking tasks to skills and quality safeguards. The OECD AI Principles call for human-centered values, transparency/explainability, robustness and accountability, and explicitly note that human intervention should be enabled where necessary.

1.3 | Why do these guidelines adopt a task-based approach?

Over many waves of technology, the same pattern keeps appearing. When a powerful general-purpose technology spreads, we see short-term job loss and disruption, followed by new jobs and higher productivity over time. Over the past half century this cycle has shifted economies toward high-skill work and expanded services and finance. That history matters, but it does not fully explain what is happening now with AI that reads, writes, summarizes, and reasons over context.

Two ways of thinking have shaped how policymakers view these shifts.



Skill-biased technical change

New technologies reward highly skilled workers and replace lower-skilled, routine work. This led many countries to focus on STEM and coding. Yet the data in the 1990s and 2000s complicated that story. Even as automation grew, many lower-skill service jobs expanded rather than disappeared.



Tasks framework

Instead of focusing on whole jobs, it looks at the tasks inside those jobs. It observes that information technologies tend to replace routine tasks in the middle of the skill range while leaving non-routine tasks at the top and bottom to grow. This helps explain “job polarization,” where employment and wages rise at the top and bottom and stall in the middle.

Generative and “contextual” AI push us to go further. These tools do not only touch routine work. They reach into parts of professional and administrative jobs that were once seen as safe. Drafting, summarizing, searching large bodies of text, and proposing first passes now take seconds. At the same time, advances in robotics combined with AI are moving this shift beyond information work into the physical world. Intelligent machines are beginning to perform warehouse operations, precision agriculture, inspection and maintenance tasks, and even aspects of healthcare delivery. This convergence means that the impact of AI will not be limited to “desk jobs” but will reshape both knowledge-based and manual occupations.

¹ <https://www.oecd.org/en/topics/sub-issues/ai-principles.html>


A task-based approach gives that precision. A task is a single action on a clear object with an obvious purpose, written in plain language. “Verify identity documents to approve an application” is a task. “Handle verification” is not. Looking at work this way lets us see exactly which parts of a role AI can safely support today, which parts might be ready soon after some data and process clean-up, and which parts must remain human led because they involve judgment, empathy, or safety.

This approach also fits Qatar’s context. Because many services are delivered in both Arabic and English, terminology and job titles can vary widely across organizations. At the same time, processes often cross ministries, state-owned companies, and private providers, where the same role title may mask very different day-to-day activities. By breaking work down into tasks written in plain, bilingual sentences, organizations create a shared language that HR, operations, IT, and legal teams can use together.

Most importantly, a task-based method links directly to the practical levers organizations control. Once tasks are written clearly, they can be grouped into two buckets using simple “Can” and “Should” tests. “Can” asks about feasibility. “Should” asks about appropriateness and risk. From there we can design human review points, set measurable outcomes, and recombine tasks into future roles that people recognize. Training plans, job profiles, and performance measures then follow from the new task mix.

To make this concrete, the guidelines that follow use one consistent example: **the Information Security Analyst**. By breaking that role into tasks, scoring each task on Can and Should, and then recombining tasks into a future role, we show how to move from description to delivery. The same method can be applied to other roles such as a receptionist, a bank operations clerk, or a field technician. The result is a disciplined way to adopt AI that protects trust, raises quality and speed, and gives workers a clear route to learn and progress.

Before diving into each step, it helps to see the overall flow of how this guideline enables responsible, effective AI-driven job transformation. The method follows a logical, human-centered progression that moves from understanding today’s work to designing tomorrow’s roles and equipping people with the right skills.



02 Scope and Application

While these guidelines are non-binding, it is strongly recommended that:



Organizations and institutions across sectors review, adopt and adapt the guidelines for their specific contexts. Although micro approaches alone will not suffice to address the broader challenges, they can provide valuable examples of implementation within a controlled environment.




Public sector adopts a whole-of-government perspective to drive AI investments and interventions to maximize public value and enhance overall living standards.

To ensure these guidelines remain practical and effective over time, MCIT encourages organizations to provide feedback on their implementation of the AI Workplace and Job Transformation Guidelines to (dipd@mcit.gov.qa). Given AI's accelerated development rate, MCIT may also conduct an annual review of the guidelines to maintain alignment with evolving global best practices, emerging industry standards, and regulatory updates.




03 Phases of AI Workforce and Job Transformation

The AI Workforce and Job Transformation approach has two phases:



Together, these phases ensure that technology adoption protects trust, maintains quality, and enhances the human contribution across Qatar's workforce. AI-enabled workplace transformation should be supported by security-by-design, proportionate governance and clear operational controls. Before introducing AI into workplace processes, organizations should consider the sensitivity of the data, the level of AI autonomy, the impact of the task, the systems the AI tool can access and the safeguards required to prevent misuse, errors, unauthorized actions or loss of accountability.

Phase 1 | Redesign the Work



Phase 2 | Empower and Upskill




Figure 1: Phases of Human-Centered AI workplace transformation

The following section defines each of the phases and provides a set of high-level guidelines. Private and public entities are encouraged to consider the following best practice guidelines across the various phases of deployment.



Phase 1 Redesign the Work 3.1				Phase 2 Empower and Upskill 3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

3.1 Phase 1 | Redesign the Work

About Phase 1: Phase 1 sets the foundation. Before buying tools or rewriting job descriptions, **we first describe the work as it is performed in Qatar's organizations at the level of tasks.** Writing tasks in plain language (action, object, purpose) gives managers, HR, IT, and frontline teams a shared view of the job that is specific enough to design change responsibly. With that view in hand, we judge each task on two questions: **Can AI help** (feasibility) and **Should AI help** (appropriateness). The answers sort work into three buckets: **Augment Now, Augment Soon, Keep Human Led.** Most importantly, make it clear where human review belongs. Phase 1 ends when each pilot role has a validated task list, a simple Can/Should score per task, named human-in-the-loop points, and a one-page "future role" blueprint that staff recognize.

Before applying AI to a role or task, organizations should **conduct a proportionate risk assessment to understand the potential impact of the use case.** This assessment should consider the nature of the task, the data involved, the level of AI autonomy, the potential impact on employees or service users, and the consequences if the AI output or action is incorrect. For agentic AI or AI-enabled tools that can call APIs, trigger workflows, update records, create tickets or perform multi-step actions across systems, the assessment should also consider autonomy-specific risks. These include unauthorized or unintended actions, excessive permissions, cascading errors across systems, difficulty reversing actions, weak auditability, unclear accountability and over-reliance on automated execution.

Where these risks are present, organizations should **define safeguards before deployment**, including human approval points, tool and API access limits, sandbox testing, whitelisting of permitted actions, rate limits, monitoring, logging, escalation triggers and rollback or recovery procedures. The level of control should increase with the potential impact and autonomy of the AI-enabled task.

3.1.1 Step 1 | Diagnose current work

Real transformation begins with a clear picture of what people actually do each day. In Qatar, where many services are delivered in Arabic and English and where processes often cross ministries and private partners, the most useful way to see the work is at the level of tasks.

Task-level analysis enables organizations to:

Identify activities ready for AI support. 	Identify activities that may change over time. 	Identify activities that should remain human-led. 	Create a common language across policy, HR, and operational teams. 
---	--	---	--

However, many entities in Qatar do not yet have processes documented down to Level 5 (task-level detail). To make the task-based approach work, organizations may need to first develop this level of granularity, drawing on international standards such as ISO/IEC process modelling guidelines or the APQC Process Classification Framework (PCF). This ensures that task definitions are consistent, comparable across organizations, and robust enough to serve as the foundation for AI augmentation and workforce redesign.

Phase 1 Redesign the Work 3.1				Phase 2 Empower and Upskill 3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

A task is easiest to work with when it is written as a simple sentence that states the action, the object, and the purpose.

“Verify identity documents to approve an application” is clearer than “handle verification.”

Most organizations already have the raw material to describe tasks properly. Job descriptions, standard operating procedures, training notes, could provide a useful resource for this step. Short interviews and brief periods of observation will fill the gaps. The goal is to create a plain list of the tasks that make up a role, not a high-level summary of responsibilities.

Consider a familiar role such as the information security analyst. In practice, the title hides many different activities. In a typical week the analyst might:

Talk to users to understand reported issues	Run short awareness sessions to reduce phishing risk	Review threat reports to spot patterns
Check access logs for anomalies	Validate encryption settings and firewall rules	Look into policy violations
Draft and update procedures that safeguard files	Assess risks and carry out agreed containment steps	Coordinate patches with infrastructure teams
Document incidents and lessons learned	Monitor ongoing data use to make sure fixes are holding.	

Written as clear task sentences, this role stops being a label and becomes a map of real work that managers and staff can discuss and improve.

Box 1 - Describing job roles through a task-based description, Information Security Analyst

After the tasks are written clearly, the work shifts from description to design. The next step is to check the list with the people who do the job. Run through the task descriptions for a given role through target interviews, workshops or roundtables. Read the sentences aloud with the team, correct the wording, add steps that were missed, and note where there are differences between shifts or locations. This short validation gives you a version of the work that staff recognize as accurate, which is essential for what comes next.

Phase 1 Redesign the Work 3.1				Phase 2 Empower and Upskill 3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

3.1.2 Step 2 | Assess AI augmentation

The second step is to judge each task on two questions.

1. Can AI support this task **now**, within the **next 6 to 18 months**, or later?
2. Should AI support this task, given the service you want to deliver, level of quality and the skills you want to build. Write a one-line reason for each judgement?

Asking these questions will allow you to determine the urgency of tasks automation and group the tasks into three buckets.

1. **Augment now** for repetitive, rules-based work that has clean inputs.
2. **Augment soon** for tasks that would benefit once data is tidied or systems are integrated.
3. **Keep human led** for tasks that rest on judgement, empathy or safety.


A systemic way to classify whether tasks can and should be AI-augmented, is by creating a scoring rubric. Give every task two scores from 1 to 3. “Can” is feasibility. “Should” is appropriateness. A score of 1 means low, 2 means medium, 3 means high.

- If both scores are 3, start now (**Augment now**).
- If both scores are at least 2 but not both 3, prepare and start soon (**Augment soon**).
- If either score is 1, **keep the task human-led** for now.

You do not average or add the two scores; you keep them separate and use the simple rule above to place the task into a bucket.



Choose one final number (1–3) for “Can” and one for “Should” using these plain-language anchors.





Phase 1 Redesign the Work 3.1				Phase 2 Empower and Upskill 3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

Description of the “Can” score

Use the “Can” score to answer a practical question: could a tool help without months of preparation?

For a task like “check access logs to detect anomalies,” the “Can” score is high only if the logs are accessible, fields are consistent, and you can retrieve them without manual exports. A score of 2 often means something small is missing, such as a single API, a data clean-up pass, or a clear rule sheet. A score of 1 signals that you should fix the basics first. This avoids promising “AI” when the real issue is scattered data or unclear rules.

Score What it means

1	Data is scattered or missing; rules are unclear; systems are not connected.
2	Some data and rules exist; there are workarounds; limited integration is possible.
3	Data and rules are clear; access and integrations are in place; the task is repetitive with steady volume.

Table 1 - Can AI Help? Feasibility


Description of the “Should” score

Use the “Should” score to check whether assistance fits your duty of care and service goals. Ask what happens if the tool is wrong and how easy it is to catch and correct. For an information security analyst, “draft incident timelines from logs” usually scores 3 on Should because a human verifies the final report and errors are easy to spot. “Approve policy exceptions” usually scores 1 because judgment, context, and accountability are central. When the team is split between 2 and 3, prefer the lower score unless you already have clear human review, strong audit trails, and a simple way to reverse mistakes.

Score What it means

1	High risk to rights, safety, or trust if wrong; heavy judgment or empathy required; high learning value if kept manual.
2	Moderate risk or judgment; some learning value; assistance is possible if a person remains firmly in control.
3	Low risk if wrong and easy to review; little need for judgment or empathy; low learning value to keep manual.

Table 2 - Should AI help? Appropriateness



You now have two numbers that tell different stories. “Can” says how ready the task is. “Should” says how safe and sensible it is to proceed. Keep them separate and do not average them. A task that is easy to do but unwise to rush should not jump the queue. Likewise, a task that would be appropriate in principle but lacks data or access should wait while you prepare the ground.






Figure 2: Classifying tasks into buckets

When scoring whether AI can and should support a task, organizations should also identify the intended level of AI autonomy. The higher the level of autonomy, the stronger the human oversight, technical controls and information security safeguards should be.






Phase 1 Redesign the Work				Phase 2 Empower and Upskill	
3.1				3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

For all three modes, organizations should define what the AI system may propose or do, where a human must intervene, who is accountable for the final outcome, and what technical and information security controls are required before deployment.

With the scoring guide in hand, it helps to see how the method works on a real job that many organizations in Qatar recognize. The Information Security Analyst role is a good test case because it mixes routine checks, documentation, and judgment-heavy decisions. By applying the “Can” score for feasibility and the “Should” score for appropriateness to each task in this role, we can show exactly how to sort work into augment now, augment soon, and keep human led. The example that follows walks through typical tasks, explains the scores in plain language, and states where human review sits, so readers can see how the rules translate into day-to-day practice and how the same approach can be reused across other roles.

ID	Task sentence	Can (1-3)	Should (1-3)	Bucket	Reason	Human review
T01	Discuss user-reported issues to understand symptoms and context.	2	1	Keep human-led	Conversation and empathy drive the outcome.	Supervisor samples notes weekly.
T02	Run awareness sessions and share guidance to reduce phishing risk.	2	2	Augment soon	AI helps with materials; delivery and questions need a person.	Content reviewed by security lead before use.
T03	Review threat-intelligence and malware reports to spot patterns.	3	2	Augment soon	Summaries and clustering help; analyst decides what matters.	Analyst approves alerts promoted to watch list.
T04	Check authentication and access logs to detect anomalies.	3	2	Augment soon	Feasible to flag unusual activity; thresholds need tuning.	Senior analyst reviews any access revocation.
T05	Validate encryption settings and firewall rules to protect systems.	3	2	Augment soon	Configuration checks are reliable; changes carry risk.	Change Advisory Board approves any rule change.
T06	Review and investigate policy violations.	3	2	Augment soon	Classification and triage can be assisted; intent needs judgment.	Investigator signs final disposition.


Table 3 - Example of Information Security Analyst





Phase 1 Redesign the Work				3.1			Phase 2 Empower and Upskill		3.2		
Step 1 Diagnose current work		Step 2 Assess AI augmentation		Step 3 Define human controls		Step 4 Recombine into future roles		Step 1 Translate to skills		Step 2 Embed learning in daily work	
ID	Task sentence	Can (1-3)	Should (1-3)	Bucket	Reason	Human review					
T07	Draft and update procedures that safeguard files and services.	3	3	Augment now	Drafting and versioning can be accelerated.	Security lead approves final text before publication.					
T08	Assess risks and execute containment or remediation steps.	2	1	Keep human-led	Safety and business impact require real-time judgment.	Incident commander records decisions and rationale.					
T09	Coordinate implementation of patches and controls with infrastructure teams.	3	2	Augment soon	Scheduling and reminders can be assisted; people decide timing.	Change manager signs off maintenance windows.					
T10	Document incidents and lessons so others can learn.	3	3	Augment now	Timelines and summaries can be drafted from logs.	Incident reviewer approves final report.					
T11	Monitor ongoing data use to ensure fixes are holding.	3	2	Augment soon	Flagging is feasible; context still matters for action.	Data owner reviews escalations before action.					

Table 3 - Example of Information Security Analyst



Phase 1 Redesign the Work 3.1				Phase 2 Empower and Upskill 3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

3.1.3 Step 3 | Define human control

The point of this step is to design human control on purpose, not by accident. Begin by writing clear decision boundaries. State what the system may propose, what only a person may decide, and who is accountable for the final outcome in each scenario. Then design the review path so it is specific, fast, and auditable.

Set review tiers that match impact:

- **For low-impact suggestions** such as drafting routine text, a frontline reviewer can approve or edit within minutes.
- **For medium-impact** steps such as routing, prioritization, or form validation, require same-day human confirmation before any adverse action is taken.
- **For high-impact** actions such as benefits decisions, safety controls, or anything that could materially affect a person's rights, require senior approval before the action takes effect. Write target time limits for each tier.

For example, routine approvals within 30 minutes during business hours, validations within the same day, and high-impact decisions within two hours once flagged. Define the triggers that send a case to a person. Triggers can include a confidence score below a threshold, missing or conflicting data, language or script changes, unusual combinations of fields, or any action that would deny or restrict access to a service. Keep the rule simple. Auto-approve only when all required conditions are met and documented. Send everything else to a human queue. For each trigger, name the role that reviews it and the maximum time allowed before a decision.



Wherever AI agents or AI-enabled tools are introduced into day-to-day work, Human Resources functions should play a defined role in monitoring quality, consistency, and people impact. This includes establishing clear quality benchmarks for AI-assisted outputs, validating that task outcomes meet agreed service standards, and regularly sampling work to detect errors, bias, or over-automation. HR, in coordination with line management and IT, should ensure that AI assistance does not quietly erode skill development, role clarity, or accountability. Periodic reviews should assess whether human review points remain effective, whether task boundaries are being respected in practice, and whether employees understand when to rely on AI suggestions and when independent judgment is required. Findings from these reviews should feed back into task design, training priorities, and performance management, ensuring that AI adoption consistently improves quality rather than merely increasing speed or volume.

Phase 1 Redesign the Work 3.1				Phase 2 Empower and Upskill 3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

Where AI is introduced into workplace processes, organizations should inform employees how AI is being used, what types of tasks it supports, what data it relies on, and how AI-assisted outputs may affect task allocation, quality review, performance discussions or role redesign.

AI should not be used for intrusive surveillance, excessive productivity tracking or fully automated performance management decisions. Where AI-generated insights are used to support performance management, they should be treated as one input only and reviewed by a responsible human manager. Employees should have a clear opportunity to understand, discuss and, where appropriate, challenge AI-supported assessments that materially affect them.

Organizations should also ensure that AI-related performance measures do not focus only on speed or volume. Quality, accuracy, employee capability, service outcomes, well-being and responsible use of AI should also be considered when evaluating the impact of AI-enabled workplace transformation.

Before introducing AI into a workplace task, organizations should confirm that:

The task has been assigned an impact level, such as low, medium or high



The system's role is clearly defined, including what it may suggest, draft, prioritize or trigger



The human decision-maker is clearly identified



The accountable owner for the final outcome is named



Review triggers are defined, including low confidence, missing data, conflicting information, unusual patterns or potential impact on rights, safety or access to services



Time limits for review and escalation are documented



Audit logs are retained for AI-assisted recommendations, human approvals, overrides and final decisions



Quality checks are performed periodically; and findings are used to update task design, training, thresholds and controls.



Box 2- Human-in the loop control checklist



Phase 1 Redesign the Work				3.1	Phase 2 Empower and Upskill	3.2
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles		Step 1 Translate to skills	Step 2 Embed learning in daily work

3.1.4 Step 4 | Recombine tasks into future roles

The fourth step is to recombine tasks into future roles that make sense. Keep the job familiar enough that people can see themselves in it but shift the balance toward higher value work. For an information security analyst this might mean less time copying logs and more time interpreting patterns and leading drills. For a citizen service advisor, it might mean less time retyping forms and more time solving exceptions and improving the knowledge base. Write a one-page role blueprint that states the purpose, the main responsibilities, the approximate task mix across automated, assisted and human led work, the skills and proficiency levels required, the key interfaces, and the measures that will define success.

Below is a clear example using the Information Security Analyst. You can copy the structure for any role.

Table 4 - Example Information Security Analyst - Future Role Blueprint

 Purpose	Protect the organization by spotting threats sooner, responding faster, and learning from every incident. Use assistance for heavy reading and first drafts. Keep people in charge of judgment, approvals, and communication.
 Main responsibilities	• Review assisted alerts and decided what to investigate. • Gather facts, interpret patterns, and decide on actions. • Execute and coordinate approved steps. • Draft and send clear updates to affected teams. • Document incidents and update procedures so the next response is better.
 Approximate task mix	• Automated tasks: collect logs, run standard checks, close low-risk tickets once a human has verified completion. • Assisted tasks: draft incident timelines, summarize long reports, cluster similar alerts, suggest likely root causes, prepare message templates. • Human-led tasks: risk judgments, policy exceptions, access revocations, final communications, post-incident reviews.



Phase 1 Redesign the Work				3.1	Phase 2 Empower and Upskill	3.2
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles		Step 1 Translate to skills	Step 2 Embed learning in daily work

To make comparisons across teams and sectors easier, use a shared vocabulary for tasks. You can also cross-reference recognized international libraries where helpful but keep the local description as the source of truth. This approach lets a customer-service task in a bank and a help-desk task in a ministry be compared fairly, supports consistent training design, and makes it easier to see when different groups are solving the same problem in slightly different ways.

The immediate benefit of viewing work at task level is clarity. The longer-term benefit is better judgment. With tasks described consistently, organizations can estimate how quickly change is likely to affect each part of a role and what kind of impact to expect. They can see where safeguards are needed, where skills must be strengthened, and where early pilots are most likely to succeed. Staff also recognize their own work in these descriptions and can take part in shaping how it evolves. That is the point of this phase. Understand the job as it is actually performed in Qatar's institutions today so that any transformation that follows is deliberate, fair, and credible.





Phase 1 Redesign the Work 3.1				Phase 2 Empower and Upskill 3.2	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

3.2 Phase 2 | Empower and Upskill

About Phase 2: Phase 2 takes the redesigned roles from Phase 1 and makes them real. The goal is simple: people should feel confident doing the new mix of automated, assisted, and human-led tasks, know where their judgment is still essential, and see how their careers can grow. This phase is not about starting from scratch, it builds directly on the task lists, Can/Should judgments, human-in-the-loop points, and role blueprints already created.

Organizations should ensure that AI-enabled transformation does not widen skills gaps between employees who are able to move into higher-value work and those whose tasks are more exposed to automation. When redesigning roles, organizations should identify employees and job families that may be most affected by automation, provide structured reskilling pathways, and connect learning plans to relevant national reskilling, digital skills and talent-development initiatives where available.

Organizations should also avoid measuring AI transformation only through speed, volume or cost reduction. AI may increase the quantity of outputs while reducing quality, accuracy, employee learning or service trust if it is not managed carefully. To reduce this risk, organizations should track quality and outcome measures alongside productivity measures, including error rates, rework, customer or beneficiary satisfaction, employee confidence, human override rates, escalation rates, and the effectiveness of human-in-the-loop controls.

3.2.1 Step 1 | Translate to skills

The first step is translating future roles into clear, practical skills. Start with the one-page role blueprint and break each responsibility down into the skills people need. For most roles, this will be a combination of:



Tool skills: knowing how to use the systems and assistants.



Data and analytical skills: reading dashboards, queries, or logs.



Domain knowledge: applying the rules that keep services safe and legal.



Human skills: explaining, supporting, or coordinating with others.

Phase 1 Redesign the Work				Phase 2 Empower and Upskill	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

Keep descriptions short, clear, and tied to real tasks rather than abstract competencies. Define what “good” looks like at three levels (basic, proficient, advanced) so performance can be assessed meaningfully. Test skills through real examples, not quizzes - for instance, an information security analyst might show proficiency by running a saved query to pull the right logs, interpreting a model output, and writing a short incident note that a non-specialist can follow.

Once skills are mapped, create a simple team skills heatmap to show strengths and gaps at a glance. In this context, a heatmap is a table or visual matrix that compares the skills required for future roles against the current proficiency levels of individuals or teams, using a simple scale such as basic, proficient and advanced. This allows managers and staff to see progress and focus learning where it is most needed.



3.1 Phase 1 Redesign the Work				3.2 Phase 2 Empower and Upskill	
Step 1 Diagnose current work	Step 2 Assess AI augmentation	Step 3 Define human controls	Step 4 Recombine into future roles	Step 1 Translate to skills	Step 2 Embed learning in daily work

3.2.2 Step 2 | Embed learning in daily work

Skills only stick if they are practiced. Instead of long training programs, design learning that mirrors the real tasks. Each activity should connect directly back to a task sentence from Phase 1 so the link is obvious. The best rhythm is short cycles of: learn a concept, practice it on a real example, then use it the same week.

Support learning with practical aids: one-page checklists, annotated screenshots, or short clips that demonstrate a standard move. Keep materials bilingual (Arabic and English) where needed. Protect small blocks of learning time in the roster so staff can practice without penalty.

Make feedback and coaching part of the flow of work. Pair less experienced staff with mentors who review their outputs, run quick calibration sessions where the team compares decisions on a few shared cases, and keep a simple log of when and why staff overrode the assistant. These small routines help refine thresholds and rules while showing staff that their feedback improves the system. For example, a weekly review of five incidents by security analysts may reveal two content fixes and one threshold adjustment, making the following week smoother.

Organizations should measure the impact of AI-enabled workplace transformation using a balanced set of indicators. Productivity measures should be linked to the redesigned task mix and should, where possible, compare performance before and after AI adoption. These may include time saved, volume handled, turnaround time and reduction in repetitive manual effort.

However, productivity should not be measured in isolation. **Organizations should also track quality, error rates, rework, service outcomes, employee confidence, skills progression, human override rates, escalation rates and the effectiveness of human-in-the-loop controls.** This helps ensure that AI increases the quality and reliability of work, rather than only increasing the quantity or speed of outputs.

Finally, link skills back to career growth. Show how the new role creates opportunities to step up, and what concrete skills are needed to bridge into adjacent roles. Update job descriptions, measures, and pay where responsibilities have clearly shifted, so redesign feels like a fair deal rather than a hidden cutback.



04 Document control

Version	Date	Amendments	Author
1.0.0	June 2026	Draft of the Guidelines	MCIT

وزارة الاتصالات وتكنولوجيا المعلومات
Ministry of Communications and Information Technology
دولة قطر • State of Qatar



Prepared by
Digital Industry Policies Department
Ministry of Communications & Information Technology
Version 1.0.0 2026

E-Mail: dipd@mcit.gov.qa
www.mcit.gov.qa