



PLAYBOOK

Design Your Next Intelligent Intervention

A decision-first playbook for data, analytics and operations leaders

INSIDE

8 plays from outcome to measurable intervention

Two worked examples: patient access and contact centre demand

A blank worksheet and pilot plan for your team

Quick quiz: how decision-ready is your operation?

Decision-first operational intelligence · 2026 edition

INTRODUCTION

The alert that arrives in time

At 9:10 on a Monday, a billing run issues higher-than-expected charges to 30,000 customers. Within minutes, the first disputes appear in chat. By 9:40, calls to the billing queue are climbing. The same data exists in both versions of what happens next. What differs is when it reaches the person who can act.

TOO LATE TO CHANGE THE OUTCOME

10:25 a.m.

"Billing queue service level is below target."

The queue has already blown out. Customers are waiting, the billing issue is still generating contacts, and the team can only absorb the damage.

IN TIME TO CHANGE THE OUTCOME

9:45 a.m.

"Billing queue demand is forecast to exceed capacity from 10:15. Likely driver: this morning's billing run."

Suggested next steps: alert the billing owner, publish an IVR message, move billing-skilled agents."

Service level is still on target. The cause can be paused at source while the floor prepares.

The difference is forty minutes, and those forty minutes decide whether customers wait, whether the problem keeps growing, and whether people spend their morning firefighting or fixing. **An alert has its greatest value while there is still time to act on it.** This playbook shows you how to design that kind of alert, around one decision at a time.

WHAT'S INSIDE

How to use this playbook	Who to involve, how to run a 90-minute session and what your team will leave with	3
Start with the decision	The decision window and the Detect → Interpret → Act → Learn loop behind every play	4
The eight plays	Outcome to success measure: purpose, questions, a worked example, pro tips and an output checklist for each	5
Worked example: patient access	A completed design protecting a four-hour initial-contact target in healthcare	13
Worked example: contact centre	Tracing a service-level drop back to an upstream campaign error, minute by minute	14
How emite supports every play	Where connected data, consistent measures, role-based views and AskEmite fit	17
Worksheet and 30-day pilot plan	A blank design canvas for your own decision and a week-by-week plan to test it	18
Quick quiz	Ten questions to benchmark how decision-ready your operation is today, with next steps	20

BEFORE YOU START

How to use this playbook

This playbook turns one recurring operational decision into a well-designed intelligent intervention: an alert that recognises a meaningful condition early, explains why it matters and recommends a permitted response to the person who can act on it.

It is built for a single working session of around 90 minutes, followed by a short pilot. Work through the eight plays in order. Each one produces a defined output, and together the outputs form a complete design you can test.

Who should be in the room

Role	What they bring
Operations owner	The outcome, the decision and the authority to act
Team leader or supervisor	How the condition shows up day to day, and what people actually do
Data and analytics lead	Where the evidence lives, how reliable it is and how it is defined
Workforce or capacity planner	Available levers: staffing, skills, schedules and their lead times
Governance, risk or privacy	Boundaries on data use, automation and approvals

What you will leave with

- ☐ One clearly defined operational outcome and decision
- ☐ A decision window, evidence list and context map
- ☐ A named owner, permitted actions and guardrails
- ☐ A draft decision-ready alert and a success measure
- ☐ A 30-day pilot plan

THE QUESTION TO START WITH

If your organisation could recognise one emerging operational problem 30 minutes earlier, with enough context to recommend a useful intervention, which problem would you choose?

THE IDEA BEHIND THE PLAYBOOK

Start with the decision, then design the alert

Most alerting starts with a metric and a threshold. This playbook starts with the decision a person needs to make, then works back to the evidence, context and recommendation that would help them make it well and in time.

Gartner predicts that **by 2027, 50% of business decisions will be augmented or automated by AI agents for decision intelligence**. Designing interventions around specific decisions is how organisations make sure that capability lands where it creates value.

THE DECISION WINDOW

The period between identifying a meaningful emerging risk and the point at which an intervention can no longer materially improve the outcome. Minutes in a contact centre, hours in field service or patient access, days in planning cycles. **Intelligence earns its value when it reaches the right person inside that window.**

THE INTERVENTION LOOP

01 Detect Recognise a meaningful change against an expected baseline	02 Interpret Explain impact, contributors and time available	03 Act Recommend a permitted, feasible response to the owner	04 Learn Measure what changed and refine the design
--	--	--	---

The eight plays at a glance

Play	Question	Output
01	Which outcome matters?	One outcome statement with a target
02	What is the decision window?	Decision, deadline and action lead times
03	What signals are available?	Evidence list with sources and freshness
04	What context is missing?	Context map: systems, baselines, constraints
05	Who owns the decision?	Owner, escalation path and delivery channel
06	Which actions are permissible?	Action list with guardrails and approvals
07	What should the alert communicate?	A draft decision-ready alert
08	How will success be tested?	Success measures and a pilot plan

PLAY

01

Choose the outcome that matters

Which service, customer or business measure do we want to protect or improve?

Purpose: anchor the design in a single outcome that is frequent, measurable and consequential, so every later choice has a clear reference point.

Questions to work through

- Which outcome do customers, patients or the business feel directly when it slips?
- Is there an agreed target, and is it measured consistently today?
- How often does the outcome come under pressure: daily, weekly, seasonally?
- What does a miss cost in customer, financial or compliance terms?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

Outcome: initial patient contact within four hours of referral.

Target: at least 90% of referrals contacted within four hours.

Why it matters: first contact sets the pace for the whole patient journey, and delays are felt directly by patients and referrers.

Also applies to: *contact centre service level for a priority queue; same-day field service completion; next-day delivery promise in retail fulfilment.*

Pro tips

- Choose an outcome that comes under pressure often enough to learn from within a 30-day pilot.
- Prefer outcomes with an existing, agreed definition; you can refine later.
- Write the outcome as a sentence a frontline team leader would recognise.

Output checklist

- ☐ One outcome statement
- ☐ A target and how it is currently measured
- ☐ A note on the business impact of a miss

Define the decision window

When is intervention most valuable, and when is it too late?

Purpose: identify the specific decision that protects the outcome, and the time available to make it, so the alert can be timed to arrive while action is still possible.

Questions to work through

- What decision, made earlier, would most often prevent a miss?
- How far ahead can a developing problem realistically be recognised?
- What is the lead time of each possible action (how long before it has an effect)?
- At what point does the decision stop being useful and escalation take over?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

Decision: rebalance scheduling capacity and prioritise at-risk referrals.

Window: roughly 45–120 minutes before referrals reach the four-hour mark.

Action lead times: reassigning trained staff about 10 minutes; opening cross-location availability about 5 minutes.

Deadline: earliest at-risk referral target time minus the action lead time.

Also applies to: *minutes for contact centre queues; hours for field dispatch and patient access; days for workforce planning and inventory.*

Pro tips

- Express the window in the operation's own rhythm: intervals, shifts or cut-off times.
- If the longest action lead time exceeds the window, look for a faster action or an earlier signal.
- Note the "too late" point explicitly; it defines when escalation replaces intervention.

Output checklist

- ☐ A one-sentence decision statement
- ☐ The decision window and deadline
- ☐ Lead times for each candidate action

Identify the signals and evidence

What operational evidence establishes that the condition is developing?

Purpose: list the specific data that would let someone recognise the condition early and trust what they are seeing.

Questions to work through

- Which leading indicators move before the outcome slips?
- Where does each signal live, and how fresh is it?
- Which signals are facts (directly measured) and which are estimates or forecasts?
- What is the minimum evidence needed to act with confidence?
- Which upstream business events (campaigns, billing runs, digital releases) drive demand into the operation, and what are their earliest signals?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

Referral intake vs comparable weekday baseline (referral management, 5-minute refresh).

Referrals within 60 minutes of target (referral management, near real time).

Scheduler capacity available vs plan (workforce management, 15-minute refresh).

Inbound demand and response times (contact centre platform, near real time).

Also applies to: queue volume, handling time and adherence in a contact centre; job backlog and technician status in field service; order volume and carrier status in fulfilment.

Pro tips

- Record the freshness and owner of every signal; both appear in the final alert.
- Favour three or four strong signals over a long list of weak ones.
- Flag any signal whose definition differs between teams; resolve it in Play 04.

Output checklist

- ☐ An evidence list with source, owner and freshness
- ☐ Facts and estimates clearly separated
- ☐ Known definition issues logged

Fill the context gaps

Which related systems, baselines or constraints change what the signals mean?

Purpose: add the context that turns a moving number into an operational judgement: what is normal, what else is happening, and what limits the response.

Questions to work through

- What does "normal" look like for this time, day, season and location?
- Which known events (holidays, campaigns, outages) should change expectations?
- Which other systems hold information the decision depends on?
- Which metric definitions need to be agreed or versioned?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

Baseline: same interval on comparable weekdays over recent weeks, adjusted for public holidays.

Related context: cross-location appointment availability (scheduling system), staff skills (workforce management).

Constraints: service eligibility rules; clinical prioritisation policies.

Definition agreed: "contacted" means a two-way contact, not an attempted call.

Also applies to: campaign calendars in a contact centre; weather and network data in utilities; carrier cut-offs and stock positions in retail.

Pro tips

- Write down one consistent definition for each key measure, with an owner. Good data culture starts here.
- Exclude past incident periods when building baselines so they do not redefine normal.
- If a key context source cannot be connected yet, note it as a pilot limitation rather than a blocker.

Output checklist

- ☐ A context map of systems and baselines
- ☐ Agreed definitions for key measures
- ☐ Known events and constraints listed

PLAY

05

Assign decision ownership

Who has the responsibility and authority to act?

Purpose: make sure the alert reaches a role that can act on it, with a clear path if they cannot.

Questions to work through

- Which role is accountable for the outcome during the shift?
- Does that role have authority for every action on the list, or do some need approval?
- How should the alert reach them: dashboard, mobile, collaboration tool?
- What happens if the alert is not acknowledged in time?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

Owner: Patient Access Manager (on shift).

Escalation: Service Delivery Manager if unacknowledged after 10 minutes.

Channel: operational dashboard plus mobile notification for high-priority conditions.

Informed: scheduling team leaders see the prioritised referral list.

Also applies to: *workforce manager for queue rebalancing; dispatch lead for field jobs; fulfilment manager for carrier changes.*

Pro tips

- Route by role, not distribution list. An alert sent to everyone is often acted on by no one.
- Different roles may need different views of the same condition; note them now.
- Agree an attention budget: how many interrupting alerts per shift is sustainable?
- Look upstream: the person best placed to fix the cause may sit outside the operation. Include them as an owner where it makes sense.

Output checklist

- ☐ Named owner role and escalation path
- ☐ Delivery channel agreed
- ☐ Supporting roles and their views noted

Define permissible actions and guardrails

What can be done, by whom, and within which boundaries?

Purpose: list the interventions the organisation can actually take, so recommendations are feasible and appropriate rather than generic.

Questions to work through

- Which actions have helped in the past when this condition developed?
- What are the preconditions and constraints for each (skills, capacity, cost)?
- Which actions can the owner take directly, and which need approval?
- Is any action suitable for automation within pre-approved limits?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

Rebalance capacity: move trained staff from administrative follow-up (owner may approve; max 60 minutes).

Cross-location booking: surface availability at other sites (owner may approve).

Prioritise at-risk referrals: order the work list using established service and clinical policies.

Guardrails: clinical prioritisation and treatment decisions remain with qualified clinical staff; patient identifiers masked in notifications.

Also applies to: *reallocating cross-skilled agents or offering callbacks; extending shifts or rescheduling jobs; switching carriers or adjusting cut-off messaging.*

Pro tips

- Write each action with its preconditions, lead time, approval level and whether it is reversible.
- Start with "recommend" rather than "automate"; automation can follow once acceptance is proven.
- Make the boundaries explicit in the alert so recipients know what is and is not being suggested.

Output checklist

- ☐ Action list with preconditions and lead times
- ☐ Approval level for each action
- ☐ Guardrails documented

Design the alert

What should the alert communicate so the owner can act without investigating first?

Purpose: assemble the outputs of Plays 01–06 into a short, decision-ready briefing.

Questions to work through

- Does the opening line state the condition in business terms?
- Is the impact clear, with the time available to act?
- Are verified facts separated from likely contributors?
- Are the recommended options permitted, feasible and specific?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

"Initial-contact performance is at risk within the next 45 minutes. Referral intake is 42% above the comparable weekday baseline and 6 of 8 schedulers are available. Eighteen referrals are approaching the four-hour target.

Recommended next steps: review cross-location capacity; consider temporary reassignment of trained staff; prioritise contacts under approved service and clinical policies."

Owner: Patient Access Manager · **Reassess:** 15 minutes · **Data as of:** 9:03 a.m.

Also applies to: *any operation: the structure stays the same; only the evidence and options change.*

Pro tips

- Lead with impact; put supporting metrics second.
- Label inferences ("likely driver") so they read differently from facts.
- Always include a review time; it turns a message into a managed intervention.

Output checklist

- ☐ Draft alert text
- ☐ Passes all seven tests (see below)
- ☐ Reviewed by the owner and a team leader

THE SEVEN TESTS OF A DECISION-READY ALERT

1 Change What happened?	2 Impact Why does it matter?	3 Evidence What is contributing?	4 Window How long to act?
5 Action What can we do?	6 Owner Who should respond?	7 Outcome Did it work?	

Define how success will be tested

How will we know the intervention made a difference?

Purpose: decide in advance what improvement looks like, how it will be measured and what would prompt a change of design.

Questions to work through

- Which outcome measure should improve, and by when?
- How will you know whether alerts were useful (acknowledged, accepted, acted on)?
- What comparison will separate the effect of the intervention from normal variation?
- What feedback will you capture from the people receiving alerts?

WORKED EXAMPLE · HEALTHCARE PATIENT ACCESS

Outcome measures: four-hour contact compliance; median time to first contact; referral backlog age.

Decision measures: time to acknowledgement; recommendation acceptance; actionable alert rate.

Comparison: pilot at two sites first; compare with matched weeks and non-pilot sites.

Review: weekly 20-minute review with the owner and data lead.

Also applies to: *service-level adherence and abandonment; first-time fix and same-day completion; on-time dispatch and delivery promise kept.*

Pro tips

- Measure decisions as well as outcomes: an unused alert cannot improve anything.
- Record when a recommendation is modified, and how; it is the richest design feedback you will get.
- Remember that a successful intervention prevents the predicted problem, so judge prediction accuracy on cases where no action was taken.

Output checklist

- ☐ Outcome and decision measures defined
- ☐ Comparison approach agreed
- ☐ Review cadence and owner set

The completed design: patient access example

Here is the full output of the eight plays for the healthcare patient-access example. Your completed worksheet should read in a similar way: specific, testable and owned.

Element	Worked example
01 Outcome	≥90% of referrals receive initial patient contact within four hours
02 Decision + window	Rebalance scheduling capacity and prioritise at-risk referrals, 45–120 minutes before targets are reached
03 Signals + evidence	Referral intake vs baseline; referrals within 60 minutes of target; scheduler capacity vs plan; inbound demand
04 Context	Weekday baselines adjusted for public holidays; cross-location availability; staff skills; eligibility rules; agreed "contacted" definition
05 Owner + escalation	Patient Access Manager; escalates to Service Delivery Manager after 10 minutes unacknowledged
06 Actions + guardrails	Rebalance capacity (max 60 minutes); cross-location booking; prioritise at-risk referrals. Clinical prioritisation remains with clinical staff; identifiers masked
07 Alert	Initial-contact performance at risk within 45 minutes; evidence; three recommended steps; owner; reassess in 15 minutes; data as of time
08 Success	Four-hour compliance, time to first contact, backlog age; acknowledgement and acceptance; two-site pilot vs matched weeks

Illustrative scenario and figures; not a customer result or live product screen.

When service level is the symptom

Tracing contact demand to its upstream source

Service level is the measure most contact centres use to judge operational health, and for good reason: customers feel it directly. It is also an outcome of the whole business. When it falls, the instinctive response is to look at the floor: answer faster, pull agents off offline work, shorten conversations.

Often, though, the cause sits somewhere else entirely. A promotional email goes out with a faulty discount code. A billing run produces unexpected charges. A mobile app release breaks a step in checkout. Each one generates a wave of contacts that lands on the contact centre minutes or hours later.

Queue monitoring shows the effect once it reaches the queue. Connecting upstream business data to interaction data shows the likely cause as it develops, which widens the decision window and brings a new kind of action into reach: **fixing the source, not only absorbing the demand.**

Common upstream drivers of contact demand

Upstream event	Early signal	Where it lives	Typical contact impact
Campaign or offer issue	Promo redemption failures; checkout errors on offer orders	Marketing automation, eCommerce platform	Order and payment enquiries, complaints
Billing run anomaly	Unusual charge values or volumes in a billing batch	Billing / ERP	Billing disputes, payment arrangements
Digital release or outage	Error rates, failed logins, abandoned journeys	Web and app analytics, monitoring	Account access and "how do I" contacts
Payment or delivery disruption	Payment declines; carrier delay events	Payments gateway, logistics systems	"Where is my order" and refund contacts
Policy, price or product change	Change go-live; spike in related page views	Change calendar, CMS, product systems	Clarification and cancellation contacts

Illustrative examples. The right drivers to monitor depend on each organisation's products, channels and systems.

SECOND WORKED EXAMPLE · CONTACT CENTRE

A morning campaign, minute by minute

Illustrative scenario. Figures and workflow show how a context-aware intervention could be designed; they are not customer results.

7:30	MARKETING	Campaign sent A promotional email reaches 400,000 customers. The discount code has been configured incorrectly.
7:45	ECOMMERCE	Checkout failures rise Failed checkouts on orders using the code climb from under 1% to 18%. No queue has moved yet.
7:55	DIGITAL CHANNELS	Early contact signals Chat and web-form contacts mentioning the offer begin to rise; IVR selections for "orders" edge above the comparable Monday.
8:05	DECISION-READY ALERT	Cause and impact connected The upstream checkout errors, the campaign send and the early contact signals are recognised as one developing condition. Service level is still within target.
8:10	ACTION	Fix the source, prepare the floor The eCommerce owner suspends and reissues the code. The contact centre adds an IVR and website message, enables callbacks and publishes a knowledge article for agents.
9:30	OUTCOME	Learn Service level held within target. Offer-related contacts, time from upstream event to fix, and callback take-up are recorded for review.

THE DECISION-READY ALERT AT 8:05

"Orders queue demand is forecast to exceed capacity from about 8:30, with service level likely to fall below target by 8:45. Checkout failures on orders using this morning's campaign code have risen to 18% since 7:45, and offer-related chat contacts are climbing. **Likely driver: campaign discount code error** (inferred from timing and contact reasons).

Recommended next steps: alert the eCommerce owner to suspend or correct the code; publish an IVR and website message; enable callbacks for order enquiries; brief agents with a short knowledge article."

Owners: Contact Centre Operations Manager (demand) · eCommerce Lead (cause) · Reassess: 8:20 · Data as of: 8:04

The completed design: contact centre example

The same eight plays, applied to upstream contact demand.

Element	Contact centre worked example
01 Outcome	Service level of 80% of calls answered within 30 seconds on the Orders queue, alongside avoidable contacts per 1,000 orders
02 Decision + window	Fix the upstream cause and prepare the floor, 30–60 minutes before the demand wave peaks in the queue
03 Signals + evidence	Checkout failure rate by offer; campaign send log; chat and web-form contact reasons; IVR selections; queue arrivals vs forecast
04 Context	Campaign and release calendar; billing run schedule; interval-of-week baselines; agreed definition of an "offer-related contact"
05 Owner + escalation	Contact Centre Operations Manager (demand) and eCommerce Lead (cause); escalates to Head of Digital if the source is not addressed within 20 minutes
06 Actions + guardrails	Suspend or correct the code (eCommerce Lead approves); IVR and web messaging; callbacks; knowledge article; cross-skilled agent moves. Customer pricing commitments honoured under existing policy
07 Alert	Demand forecast to exceed capacity from 8:30; likely driver identified and labelled as inferred; next steps for both owners; reassess at 8:20
08 Success	Service level held; offer-related contacts avoided; time from upstream event to fix; compared with previous campaign incidents

Illustrative scenario and figures; not a customer result or live product screen.

THE INSIGHT

The highest-value action in this scenario sits **outside the contact centre**. A decision-ready alert can route to the person who owns the cause as well as the person who owns the queue, so the business fixes the problem at source while the floor prepares for the contacts already on their way.

Questions to take back to your team

- Which upstream events have caused your last three service-level drops?
- How long after each event did the impact become visible in the queue?
- Who outside the contact centre could have acted, and did they know in time?

PUTTING THE PLAYBOOK TO WORK

How emite supports every play

Each play depends on connected, consistent data and on getting the right insight to the right person. This is where emite fits.

Play	What the play needs	How emite helps
01 Outcome	A trusted, consistently measured outcome	Advanced Analytics applies consistent business measures, so the outcome means the same thing across teams and systems
02 Window	Visibility early enough to act	Advanced iPaaS brings in upstream business data alongside interaction data, so developing conditions can be seen as they emerge, often before the impact reaches the queue
03 Signals	Evidence from many systems in one place	Advanced iPaaS connects contact centre, workforce, CRM, digital, billing and operational systems into a single, governed data foundation
04 Context	Baselines, calendars and agreed definitions	Advanced Analytics provides baselines and consistent definitions; business calendars and events can be brought in as context
05 Owner	The right view for each role, including upstream owners	Advanced Visualisation delivers role-based views, from frontline team leaders to executives and business owners outside the contact centre
06 Actions	Options grounded in current capacity and context	AskEmite offers suggested next steps grounded in the connected data, with people approving consequential decisions
07 Alert	A briefing that explains what changed and why	AskEmite provides anomaly awareness and contextual explanations, and lets people ask follow-up questions in natural language
08 Success	Outcome tracking over time	Advanced Analytics and Advanced Visualisation track outcome and decision measures, so each intervention can be reviewed and refined

IN THE CONTACT CENTRE SCENARIO

emite connects the campaign, eCommerce and digital-channel data to the interaction data, so the rise in checkout failures and the first offer-related contacts are visible alongside queue forecasts. AskEmite surfaces the anomaly, explains the likely driver and suggests next steps, giving both the contact centre and the eCommerce owner the foresight to act while service level is still within target. **Your data answers back.**

Design your intervention

Use one worksheet per decision. Fill each box with the outputs from the matching play.

01 OUTCOME <i>What result do you want to protect or improve, and what is the target?</i>	02 DECISION + WINDOW <i>Which decision matters, how long is there to act, and what are the action lead times?</i>
03 SIGNALS + EVIDENCE <i>Which data establishes the condition? Source, owner, freshness.</i>	04 CONTEXT <i>Baselines, known events, related systems, agreed definitions, constraints.</i>
05 OWNER + ESCALATION <i>Who receives the alert, who decides, and who is next if there is no response?</i>	06 ACTIONS + GUARDRAILS <i>What is permitted and feasible? What needs approval? What stays with people?</i>
07 ALERT DRAFT <i>What changed · why it matters · next step · owner · when to reassess</i>	08 SUCCESS + FEEDBACK <i>Outcome and decision measures, comparison approach, review cadence.</i>

Review the completed design with operations, data, security and governance owners before deployment.

FROM WORKSHOP TO PILOT

Your 30-day pilot plan

A focused pilot proves value quickly and builds the evidence needed to scale. Keep the scope to one decision, one or two sites or teams, and a fixed review rhythm.

Week	Focus	Activities	Done when
1	Prepare	Connect evidence sources; confirm definitions; set baselines; brief the owner and team leaders	Evidence refreshes within agreed freshness; owner signed off the alert format
2	Run in shadow	Generate alerts without notifying; compare with what actually happened	Alert timing and volume look sensible; thresholds tuned
3	Go live	Deliver alerts to the owner; capture acknowledgement, decision and action	Every alert has a recorded decision
4	Review	Compare outcomes with matched weeks or non-pilot sites; gather owner feedback	Decision to scale, refine or stop, with evidence

Signals your pilot is working

- ☐ Owners acknowledge most alerts within the agreed time
- ☐ Recommendations are accepted, or modified in consistent and useful ways
- ☐ Interrupting alerts per shift stay within the attention budget
- ☐ The outcome measure improves against the comparison
- ☐ Team leaders say the alert saves them investigation time

QUICK QUIZ · 5 MINUTES

How decision-ready is your operation?

Answer each question for the operation you know best. Tick one option per question, then add up your points.

1. How are your most important operational outcomes defined?

☐ **A** Informally, or differently by team

0 points

☐ **B** Defined, but not linked to real-time monitoring

1 point

☐ **C** Defined, agreed and monitored in near real time

2 points

2. How are most of your alerts triggered?

☐ **A** Fixed thresholds on single metrics

0 points

☐ **B** Thresholds with some time-of-day adjustment

1 point

☐ **C** Contextual, calendar-aware baselines

2 points

3. When a problem develops, how many notifications do people receive?

☐ **A** One per metric that moves

0 points

☐ **B** Some grouping, but often several

1 point

☐ **C** One alert per condition, with evidence attached

2 points

4. Do your alerts say how much time is available to act?

☐ **A** No

0 points

☐ **B** Sometimes, informally

1 point

☐ **C** Yes, with a decision deadline or review time

2 points

5. Who receives your alerts?

☐ **A** Distribution lists or everyone

0 points

☐ **B** A team, without a clear owner

1 point

☐ **C** A named owner role with authority and escalation

2 points

6. Do alerts include recommended actions?

☐ **A** No, people decide from scratch

0 points

☐ **B** Generic guidance

1 point

☐ **C** Specific options that are permitted and feasible

2 points

7. Can you combine contact centre, workforce and operational data in near real time?

☐ **A** No, each lives in its own tool

0 points

☐ **B** Partly, with manual effort

1 point

☐ **C** Yes, through governed integration

2 points

8. Are key metric definitions consistent across teams and systems?

☐ **A** Rarely

0 points

☐ **B** For some measures

1 point

☐ **C** Yes, agreed, owned and versioned

2 points

9. Do you record what happened after each alert?

☐ **A** No

0 points

☐ **B** Occasionally, or in free text

1 point

☐ **C** Yes: acknowledgement, decision, action and outcome

2 points

10. How is AI-supported insight governed in your operation?

☐ **A** No defined approach yet

0 points

☐ **B** Some policies, applied inconsistently

1 point

☐ **C** Facts and inferences labelled; approval tiers defined

2 points

YOUR SCORE: _____ / 20

QUIZ RESULTS

What your score means

0–7 Getting started

Alerting is largely threshold-based and disconnected from decisions. That is common, and it means early improvements can deliver quick, visible value.

Recommended next steps

- Run Plays 01–05 for one high-value decision this month.
- Agree consistent definitions for the three or four measures that matter most.
- Read: **Why the Next Generation of Alerts Must Be More Intelligent.**

8–14 Building momentum

You have good foundations, with gaps in context, ownership or follow-through. The biggest gains usually come from connecting data and closing the loop.

Recommended next steps

- Complete all eight plays and run the 30-day pilot.
- Group related signals into one condition per decision, with a named owner.
- Read: **Anatomy of an Alert Worth Acting On** and **From AI Insight to Meaningful Recommendations.**

15–20 Decision-ready

Your alerting is designed around decisions and governed well. The opportunity now is scale and continuous learning.

Recommended next steps

- Extend the approach to further decisions and teams using the worksheet.
- Formalise decision contracts, prioritisation and outcome evaluation.
- Read: **Engineering Decision-Ready Intelligence** (technical deep dive).

HOW EMITE HELPS

From connected data to timely action

emite brings contact centre and enterprise data together through **Advanced iPaaS**, applies consistent business measures and baselines through **Advanced Analytics**, and delivers role-based operational views through **Advanced Visualisation**.

AskEmite builds on that foundation with anomaly awareness, contextual explanations and suggested next steps, alongside natural-language interrogation, so the people who own each decision can see what changed, understand why it matters and explore what to do next. Your data answers back.

CONTINUE READING

- › Why the Next Generation of Alerts Must Be More Intelligent
- › Healthcare Patient Access: Turning Operational Signals into Timely Intervention
- › Anatomy of an Alert Worth Acting On
- › From AI Insight to Meaningful Recommendations
- › Engineering Decision-Ready Intelligence

Book a personalised demo ›

emite.com



Illustrative scenarios and figures are examples, not customer results or live product screens. This playbook does not constitute legal, clinical or regulatory advice. Gartner statistic: "Gartner Announces the Top Data & Analytics Predictions", 17 June 2025. © 2026 Prophecy International Holdings Ltd.