



DIAGNOSIS BEFORE PRESCRIPTION:

*A PROBLEM-LED METHOD FOR TECHNOLOGY
ADOPTION IN THE HOSPITALITY INDUSTRY*

Assoc Prof Dr Michael A. Kruesi &
Assoc Prof Dr Kiattipoom Kiatkawsin



EXECUTIVE SUMMARY

Hotel commercial teams sometimes fall into the trap of adopting the latest technology, especially the latest Artificial Intelligence (AI) tools, based on speculation or a fear of missing out, without taking the time to thoroughly diagnose the problems those tools are meant to solve. Broad AI hype and heavy vendor marketing reinforce the supposed urgency, and many operators are convinced that trying something to keep up with the latest trend is better than doing nothing. However, the costs of this approach may outweigh the return, suggesting that a more measured AI adoption strategy is worth considering. This paper argues for reversing the invest-first, justify-later sequence of AI integration in favour of a more deliberate and conservative strategy that begins with taking a step back and evaluating whether the tasks performed each day would genuinely be enhanced through automation. Drawing on a workshop conducted with senior commercial leaders at the 2026 HSMAI Commercial Strategy Conference in Singapore, we present a diagnosis-first, prescription-later mindset. Specifically, we suggest a method where teams begin with a needs assessment, mapping their typical working day by classifying every task into one of three categories, before passing the problematic tasks through an assessment framework we call the 4D filter. When applied during the HSMAI workshop, the method showed that roughly a third of the tasks identified by participants were not worth the effort of automation and therefore not worth investing in. We suggest that running this exercise with your team before adopting any AI tools is time well spent, and the paper closes with a practical guide for doing so.

INTRODUCTION

Commercial leaders in hospitality are currently making technology decisions under pressure that often has little to do with their own daily operations and, most importantly, sits outside their expertise. Every week seems to bring another AI-powered tool promising greater effectiveness and efficiency, and there is increasing evidence that adoption is growing rapidly, with little coverage of how these implementations worked out in the long run. As new tools and reporting about them proliferate, the pressure to adopt is not to be underestimated and creates a fear of being left behind. As such, adoption decisions can begin without a thoroughly diagnosed problem and instead be driven by reactive responses.

We understand this pressure and do not dismiss it. Technology, notably AI, is moving fast, the tools are improving rapidly, competitors are genuinely moving, and standing still carries risks of its own. It also cannot be denied that the industry has been punished before for adopting technology too slowly. However, a strategy based on fear, whether of missing out or of falling behind, is a poor basis for long-term success, and acting on it impulsively can result in costly mistakes.

This paper therefore argues that the sequence of technology adoption must be diligently planned rather than driven by external forces such as market hype. Before any commercial team evaluates a tool, it should be able to name, precisely, the problems that the tool would solve, where those problems sit in the team's daily workflow, and what they cost in time and energy. Diagnosis can be tedious and even overwhelming, which is why it is often skipped. However, it provides a more reliable foundation for an adoption decision, and the hours spent on it can save an organisation considerable time and money in the long run. The argument laid out in this paper is based on extant literature, general market trends and our own workshop observations, informed by consultation with commercial hospitality managers from across the Asia Pacific, and the sections that follow present the reasoning, the method, and a practical guide for conducting the same diagnosis with your own team.

WHY AI ADOPTION OFTEN FAILS

THE TECHNOLOGICAL ADOPTION CYCLE

Technology adoption in many hospitality organisations follows a recognisable pattern. A new category of tool emerges, early testimonies of incredible transformation circulate, and before long companies can buy in without truly understanding what the tool is for or whether it really fits their operation. The hope, perhaps, is that measurable impact will reveal itself after implementation, on the assumption that success elsewhere indicates an objectively sound strategy. Sometimes that impact does appear, when the tool is used to its full capability, but in other cases it is used at a fraction of it, by staff who do not fully understand it or are not optimally trained to use it. By the time the subscription comes up for renewal and the results have proved to be less than expected, it can be easy to simply write off the spend, with little probing into why it did not yield the expected impact. This is a pattern with some notable examples. Henn-na Hotel in Japan, which opened in 2015, chose to invest in 243 service robots to automate its workflows and save on staff costs. Just four years later, however, the hotel removed more than half of these robots after they created more work for the human staff than they were intended to save. Winding Tree is another example, having spent seven years building a blockchain travel marketplace after a token sale and industry pilots. However, after limited adoption the project was wound down in 2024, suggesting that the marketplace had not achieved the level of adoption needed to remain viable. Finally, Expedia launched a generative booking assistant in 2024 to boost sales, but the initiative was later retired after travellers discovered that the rates, availability and reservations were not actually live.

TECHNOLOGY-FIRST THINKING

The same cycle can persist because of how new technology captures managerial attention before practical implementation is fully considered. There is no doubt that these tools are increasingly capable and promising on paper, but a well-run demonstration draws the focus towards what the tool promises rather than what the business and the team actually need. The assessment of the tool's usefulness becomes a matter of what the technology could do rather than whether it is a good fit for the organisation or whether the team requires it to improve its operations. This is the distinction between a must-have and a good-to-have, or better yet, a good-to-be-seen-to-have. As a result, the decision to purchase such a technology enhancement may effectively be made before a defined need has even been identified. Unless a diagnosed problem is already on the table, adopting a new tool can even force a team to reorganise its workflow to fit the tool's capabilities, potentially leaving it worse off than before rather than having the technology fit the operation.

THE HUMAN FACTOR

Even a well-chosen tool still requires active adoption by the team, and teams do not always absorb change seamlessly or willingly. Implementing a new tool generally requires people to alter how they carry out their tasks, which can lead to frustration and inefficiency during the learning curve. Importantly, if this change is required without notice or a prior explanation of how it is expected to affect workflows and staff more broadly, it will likely be met with resistance. Employees who might otherwise welcome change and contribute constructive feedback may instead develop resentment and push back when the rollout has not considered them. Teams should therefore be consulted whenever a tool that requires changes to their workflow is being considered and told what it means for them in both the short and long term. When adoption is slow or fails entirely, the blame can be directed at the technology rather than at the manner of its introduction. This can leave an organisation with a team reluctant to adopt new technology, even if it is of real value. Ultimately, human capacity to integrate new technology is finite. Like your budget, it requires careful planning and should not be overdrawn.

ADOPTION TIMING

Adopting the right tool can still be the wrong decision if it arrives at the wrong time. AI development currently moves faster than organisations can adopt it, which creates a potential trap for early movers. An improved version may appear before the current one is even fully implemented, a likely scenario given that there is rarely any assurance of finality with such technologies. The company then faces a choice between defending its existing investment or cutting its losses and reinvesting in the upgrade. As such, the cost of being strategically late can be lower than the cost of being repeatedly early, and companies that adopt a new technology before it stabilises are more susceptible to being caught in a repeated upgrade cycle.

THE WORKSHOP AS A METHOD

SETTING AND PARTICIPANTS

The method we present here was developed for the HSMAI Commercial Strategy Conference in Singapore, held on the 15th of May 2026. We ran it as a hands-on workshop with senior commercial leaders from hotel groups across the region. Participants worked in groups through a fixed sequence of exercises, recording their contributions on flip charts and sticky notes as they went. Those contributions, along with our secondary research, inform the remainder of this paper. The workshop was designed to help the people in the room run such workshops with their own teams rather than to generate research data, and we present its outputs as an illustration of what the method surfaces rather than as formally analysed findings. We do not attribute any contribution to an individual or organisation. However, the broader patterns that emerged are nonetheless instructive because they show how experienced practitioners describe their own working day when directly asked to do so. These exercises and patterns are further outlined below.



THE DRAIN GAME

The session opened without any mention of technology adoption, so as not to lead the participants. They were asked first to think about their typical workday, then about the different tasks that make up the day, and then to write down the tasks they considered timewasters, listing each one separately on a piece of paper before any discussion with their group members took place. The exercise is deliberately simple and focused. Surprisingly, we observed that professionals can struggle to name what constitutes a timewaster, perhaps because they are rarely asked to think about it. Within minutes, however, the groups had compiled a list of various timewasting tasks, stemming entirely from each participant's own professional experience.

WORKFLOW MAPPING

In the next step, participants were asked to categorise each task using one of three sticky note colours. Green sticky notes were for high-value work that only they could do. Yellow sticky notes were for tasks that were necessary but draining, i.e. tasks they needed to complete but that were repetitive or took too long. Lastly, red sticky notes were for tasks that were unnecessary and should not exist at all, often legacy tasks carried on simply because they had always been done, or work that had little impact on the final outcomes. At this stage, what had started as a loose collection of complaints was organised into a structured diagnosis. Each frustration was categorised, and the notes together showed where effort in the workday was going and how each task was perceived in terms of importance by the person performing it. This allowed for a more task-based analysis of which parts of the workday created tangible value and which carried an opportunity cost because the time could be better spent. This is further elaborated on in the following section.



THE 4D FILTER

However, simply identifying and classifying the timewasters is only part of the picture, so the next stage required participants to make decisions about what these classifications meant in practice. To move towards an actionable prescription, every yellow (necessary but draining) and red (unnecessary) task was passed through a filter offering four strategies: 1) Delete the task altogether, 2) Delegate it to someone else, 3) Digitize it, which includes automating it, and 4) Diminish it, by reducing its frequency or being more selective about when it genuinely needs doing. We call this the 4D filter, and its purpose is to adapt task-triage schemes, which are long established in management literature, to the specific question of technology adoption. No task was exempt, and each was passed through the filter and assigned one of the four strategies. The goal of this step is to optimise the workday by freeing up time for the tasks where a team member brings the greatest value. From a technology adoption perspective, only tasks marked 'Digitize' warrant consideration of implementing a new tool.



TOOL MATCHING

After laying this groundwork, the exercise finally moved to the question of what new technology made sense to consider implementing. Rather than aimlessly browsing for the latest hospitality technology or being influenced by effective marketing, participants had a much clearer strategy based on the problem to solve and the features a tool would need to address it. This matters in a market where websites such as toolify.ai index, as of 2026, over 30,000 AI-powered tools, and searching with a defined problem leaves far less chance of being swayed by a newly released product with features you may not need. Even so, before choosing a specialised tool, we suggest testing whether a general-purpose large language model (LLM), such as ChatGPT, Claude or Gemini, can solve the problem at a comparable standard. A diagnosed problem can be tested against an LLM in an afternoon, revealing how much of the task is genuinely automatable and providing what is effectively a free pilot; if it solves most of the problem, a new specialised subscription may be difficult to justify. At the end of the session, participants left with tools matched to the tasks identified for digitisation and, more importantly, with a clearly named problem statement for each, which provided a clearer understanding of why a new tool was being considered and equipped operators to ask the right questions of potential vendors.

WHAT THE WORKSHOP SUGGESTS

AN EVEN SPLIT

When the colour-coded maps were completed, the most striking result was the distribution of the tasks itself. Across the groups, participants' days were divided roughly evenly among the three categories of high-value work, draining work, and work that should not exist. Although this is not a definitive finding, it is worth noting that, by their own assessment, participants indicated that close to a third of their working time generated little to no value or impact that they could explicitly name. This perceived wasted time is already worth further investigation before any conversation about tools to enhance the workday is even considered.

HIGH-VALUE TASKS

The tasks categorised as green, meaning high value and not readily outsourced, also highlighted clear trends worth mentioning. First, many of these tasks were based on personal relationships built between participants and their counterparts over time. Second, a further set of green tasks involved problem-solving, where each solution was slightly different from the last and not easily standardised. Tasks of this kind resist substitution by technology, and even by another person, because the value sits in the specific relationship or in the accumulated judgement of the individual. In these instances, any time saved by adopting AI may be disproportionately outweighed by the lost leverage of personal relationships and a reduction in the quality of service delivery, respectively.

DRAINING TASKS

The yellow tasks were the most contested and drew the most emotion during discussions. These are tasks the business genuinely requires, yet they are repetitive, time-consuming and generally as unchallenging as they are unrewarding. Participants described them as bringing little satisfaction and little recognition to the team members performing them. Their necessity can make them particularly frustrating because they cannot simply be dropped, leaving employees to spend scarce hours on work that matters to the operation but may contribute little to their sense of accomplishment or, more broadly, to advancing the overall performance of the team. This makes them a natural target zone for digitisation, and the case for acting on them can rest on human wellbeing as much as organisational success. Every hour recovered from the yellow zone can be redirected towards the work where the individual contributes most and is more intrinsically motivated.

UNNECESSARY TASKS

The red tasks highlighted a more concerning pattern associated with organisational legacy and corporate culture. These tasks appeared to persist largely because they had been 'grandfathered in' and had not been questioned for some time. One participant described a weekly meeting in which the team lead and management convened in person so that each could read aloud a report already prepared, circulated and checked, a ritual that a shared dashboard could potentially replace. Others described the routine double-checking of reports that a properly configured system could validate automatically, with human review reserved for key exceptions. The red category demonstrates that diagnosis can already create value before any potential purchase of new technology is considered. Deletion requires no software, no subscription and no implementation plan, only the willingness to ask what a task is truly for and whether it merits retaining its position on the daily agenda.

WHAT THIS MEANS

Taken together, the step-by-step strategy outlined above highlights a critical consideration. The exercise suggests that a meaningful share of participants' workdays was spent on tasks that either drained the person doing them or served no purpose they could clearly name, and identifying which was which required little more than a structured period of reflection. Just as importantly, the exercise clarified where automation may be less appropriate, since many of the highest-value tasks resisted substitution as they depended on relationships and human judgement that could not be easily outsourced, let alone automated. The diagnosis also identified tasks that participants believed could be deleted immediately, at no additional cost, and reduced the number of problems for which a technological solution even needed to be considered in the first place. This exemplifies how a clearer, more objective foundation can be reached for deciding where further investment of time or money is genuinely justified. Importantly, the workshop ended with a defined shortlist of problems genuinely worth addressing through augmentation or technology.

RUNNING THE METHOD YOURSELF

The pressure to adopt the latest shiny new tools will likely remain, and this paper's emphasis is not that you should ignore current trends, but that you should understand your weaknesses and whether a new technological tool would deliver a significant improvement before you buy into it. Over the coming years, long-term competitive advantage may depend less on buying in earliest and more on understanding the operation well enough to ensure that each purchase addresses a clearly defined issue. That understanding can be developed by any team willing to spend a few hours together diagnosing its problems carefully.

THE DIAGNOSIS-FIRST SESSION: FIVE STEPS

Below we set out a five-step framework for running this workshop with your team. However, feel free to adjust it to fit your context, including the timing of each step or the materials used to facilitate it.

- 1. The Drain Game (15 minutes):** On a piece of paper, participants are asked to write down the usual tasks that make up their workday which they consider to be a waste of their time. At this stage it is an individual exercise, and there should be no discussion. The silence matters here, as it keeps the more dominant voices from steering the room early, and it allows the genuine overlap between individual diagnoses to surface before any further instructions are given about the classification steps that follow. Also, do not yet mention technology or AI, as it may skew the responses.
- 2. Map the Day (30 minutes):** At this stage, participants need to classify the tasks by sorting them into one of three colour categories, green for high-value work only they can do, yellow for necessary but draining tasks, and red for low-value, time-wasting tasks with no clear output. Encourage everyone to be honest, and put your team at ease so they can speak freely. This is a brainstorming exercise, not a performance review, and a safe, open voice climate will help produce the best result.
- 3. Read the Distribution (15 minutes):** Step back and look at the colour-coded tasks across the whole team. Discuss what share of the collective day sits in yellow and red. If there are outliers or disagreements, ask the individuals to explain their position. This will help to better visualise the commonly agreed time-wasting tasks across the team and to identify patterns.
- 4. Run the 4D Filter (30 minutes):** Here the exercise moves from identifying the problems to fixing them. Take all the yellow and red coded tasks and assign each one a prescription, Delete, Delegate, Digitize or Diminish. No task may remain unassigned, because an undecided task can quietly return to being a tedious, low-value time waster. Act on the Deletes as soon as practical, as they require no new technology or budget. Take note of the tasks categorised as Delegate and Diminish for further evaluation later, and then move to the next step to assess whether a new tool is worth adopting by focusing on the Digitize category.
- 5. Prompt First (30 minutes):** For each task classified as Digitize, test the problem against a general-purpose LLM before considering any specialised tool. Give it a real instance of the task, using an approved platform and appropriately protected information, and spend time testing it. This is time well spent relative to the cost it incurs. Where an LLM can solve the problem to the required standard, the need for an additional specialised subscription becomes harder to justify. Whatever tasks remain become your shortlist for considering a new tool. Besides being a more targeted approach, this also helps to develop a clearer brief for potential vendors to find that specific tool for you, or for you to find it yourself.

Finally, we suggest repeating the session periodically, on a quarterly or annual basis. New issues will surface and new tools will be developed. There will continue to be new and exciting technologies that could have a major positive impact on your operation, but remember the golden rule:

Diagnosis Comes Before Prescription.

AUTHOR BIOGRAPHIES

DR. MICHAEL A. KRUESI

Dr. Michael A. Kruesi is an Associate Professor in the Hospitality and Tourism Management Programme at the Singapore Institute of Technology (SIT). He began his career in hospitality and tourism in 2002, working in hotels in New Zealand, Switzerland and Peru, before completing his PhD at Auckland University of Technology (AUT) in 2015. His academic career has since taken him from AUT to the National Kaohsiung University of Hospitality and Tourism in Taiwan, and now to SIT, where he has been a faculty member for seven years.

Dr. Kruesi's research and teaching are best summarised by the three pillars of sustainability: environmental protection, economic viability and social wellbeing. Specifically, his research focuses on non-equity governance structures and technology augmentation in hospitality, on organisational behaviour and leadership practices that promote positive employee outcomes, and on the measurement and benchmarking of the environmental impacts of hotels, attractions and events, together with the revenue optimisation and economic incentives of sustainability-based strategies more broadly.



DR. KIATTIPOOM KIATKAWSIN

Kiattipoom Kiatkawsin, Ph.D. is an Associate Professor and head of the hospitality and tourism division at Singapore Institute of Technology (SIT). Dr. Kiatkawsin's hospitality career started in a fine-dining restaurant in Germany before transitioning to event management in Switzerland. After getting a taste of teaching in Switzerland, he moved to Seoul, South Korea to pursue his Ph.D. at Sejong University. His academic career subsequently started in Korea before taking up the current role at SIT.

Dr. Kiatkawsin's expertise is in consumer behaviour, specifically customer experiential design, luxury consumption and technology-induced service delivery design. He has also been involved in various tourism development consultancy projects with organizations such as the UN Tourism, UN Development Programme, and the Korea Tourism Organization. He regularly contributes to the scientific community and has over 30 publications in SSCI-indexed hospitality and tourism journals. Dr. Kiatkawsin represents Singapore on the board of the Asia Pacific Tourism Association.

