CALIFORNIA
Office of Data and Innovation

GenAI innovation playbook

January, 2025

State of California
Office of Data and Innovation
401 I Street Suite 200, Sacramento, California 95814

Dear California state colleagues,

The Office of Data and Innovation (ODI) collaborates with state agencies and uses data, human-centered design, training, and technology to improve services. We envision a state government that is modern, effective, and provides exceptional service to all Californians.

Since the fall of 2023, the Office of Data and Innovation (ODI) has partnered with the Government Operations Agency (GovOps), the Department of Technology (CDT), the Department of General Services (DGS), the Department of Human Resources (CalHR), and many others to meet the directives in Governor Newsom's Executive Order N-12-23 on generative artificial intelligence (GenAI).

ODI is proud to have supported all five GenAI proof of concept (POC) teams as they partnered with innovators to build GenAI solutions in a sandbox. We:

- Provided product management support, ensuring that experts from both the business and IT divisions in a department helped guide the GenAI solution's build and testing.
- Trained leads in each department in the fundamentals of user experience research and analysis, and partnered with these leads to develop staff interview questions. These interviews helped confirm the business value delivered by each GenAI solution.
- Supported departments in their data access, governance, and planning needs as they prepared for potential production GenAI solutions.

The worksheets in this playbook were written with GenAI procurements in mind, but the practices here can be applied to non-GenAI projects as well. Doing so will improve your problem definition, reduce project risk, and improve Californians' customer experience of your programs. Applying small, intentional process changes like these at the beginning of each project will move us all toward statewide transformation in service delivery!

If you have questions or feedback about this playbook, please email us at info@innovation.ca.gov. For additional resources for improved service delivery, including accessibility and human-centered design, visit our Innovation Hub.

Sincerely,

Jeffery Marino, Director
Office of Data and Innovation

Table of Contents

Data readiness note sheet

It's important to recognize the critical importance of the underlying data for any AI solution. Comprehensive and accurate data is essential for successful AI solutions because it ensures the models learn correctly and give reliable output. The "garbage in, garbage out" mantra is still highly relevant in the GenAI domain.

In this early stage of the process, you don't need to iron out every last detail of the underlying data you'll use. However, it is very important to get an initial understanding of your data readiness. An early "gut check" on your data may shine a light on barriers that will need to be addressed or could even change your problem statement, depending on what data is available to you.

Keep in mind that for the proof of concept phase, you will only be using public data. But if a proof of concept is successful and your department moves forward to contract for a production tool, you may use more proprietary datasets.

Think about what data you will need to address your problem statement, and begin answering the questions below to conduct an initial data inventory, quality, and governance assessment.

Inventory

What is it? *List and describe the dataset(s) that will help address your challenge. Provide a simple description for each.*	
Where is it? *GenAI will rely on large, well-organized datasets to function well.* *List the on-prem databases, cloud databases, local storage, etc. where*	

the data lives.	
How much is there? *If you don't have much data or information, the models may not have what they need to perform well. It's also important to note if you have a lot of data, to better understand the compute that will be needed.* *Describe the amount of data stored (MB, GB, etc.) and the span of time/breadth.*	
What format is it in? *Many GenAI tools will require the data to be in a specific format.* *Describe the way the data is structured and stored (e.g., JSON, XML, CSV, XLSX).*	
Is it accessible? *Describe any barriers or restrictions that make it hard to retrieve or use the data. This may include subscription costs.*	

Quality

Is it very sparse? *If there is wide-spread missingness or incompleteness in your data, that can make it difficult for AI models to perform well.*	

Describe any known gaps or missing values in your datasets.	
Is it mostly consistent? *If there are large inconsistencies in how the data has been reported or processed, this can lead to unreliable outputs in the AI solution.* *Note whether the data collection and reporting methods over time are consistent. Did a change in technology change the data at any point?*	
Is it from reliable sources? *If the data is not from credible sources, it can be difficult for the AI solution to make informed decisions.* *Is the data from credible sources?*	

Governance

Who owns it? *Describe who has the legal rights and control over the data.*	
Is it scrubbed of PII? *There are strict rules when it comes to the data that can be used in the proof of concept sandbox and there cannot be any personally identifiable information. If the data will need to be scrubbed of PII, that will be an additional step in the process that should be flagged early on.* *Have you followed privacy laws and*	

guidelines to protect people's personal information?	
What steps will need to be taken to prepare the data? *Often, the data is not completely ready for prime time and will need to undergo significant cleaning, reformatting and/or migration efforts to be deemed usable. Meeting with data subject matter experts will be critical in understanding how much preparation will need to be done to get your data ready to be used.* *Will the data need to be cleaned, reformatted, or migrated to a new system to be useful?*	

Framing a problem statement for a solicitation

The state has used the Request for Innovative Ideas (RFI2) procurement vehicle for GenAI proofs of concept thus far because it allows for building and testing in a low-risk way, and creates the opportunity for competition between innovators (vendors). The RFI2 procurement vehicle is a powerful innovation, but innovator bids will only be as strong as the problem statement they're responding to. A clear problem statement is critical to ensure any solution meets the needs of your business, workforce, and the Californians you serve.

The strongest problem statements center people and their needs, **not** a tool or solution. Rather than jumping to a solution, a strong problem statement begins with a series of questions that get at who your many stakeholders and customers are, and how a solution could make their daily work and lives easier. In this way, the process ensures that any business process change or procurement ultimately meets the needs of its intended audience.

Start with a guiding question

We often hear "needs" or complaints from departments that something is broken. These are a great place to start. But expanding from a narrow, need-based question to a broader, open-ended question can spur even more interesting ideas–and, ultimately, bids.

Not that… (need)	***This… (question)***
We need to make better predictions.	How might we build better planning processes with data?
Our response teams are too slow.	How do we define response success?
These software tools are causing more problems than they solve.	What does an ideal workflow look like?

A guiding question is the central, exploratory, **human-centered question** that you and your team are trying to answer. Guiding questions are "How might we…" and "What would it look like if…" questions. These questions are expansive and intentionally general. They will need to be broken into smaller pieces to tackle. For example, the following guiding question might have multiple workstreams, pilots, and solutions: "How might we improve department service delivery for all Californians if we used GenAI to streamline our internal business processes?" The department challenges your team took to the September GenAI Innovator Showcase were guiding, human-centered questions like this.

What is a problem statement?

Once you have your guiding question and have conducted market research on it (such as desk research or an Innovator Showcase), you can frame it into a problem statement. A problem statement is a clear, concise, and actionable narrative that will be included in an RFI2 solicitation package. The problem statement describes the context of the situation and the core issue(s) you and your team are trying to solve. It is an alignment mechanism to ensure all parties understand the challenge at hand, and can move forward together towards a common goal and solution. Writing one is a process and should happen iteratively with all relevant stakeholders involved.

Problem identification

To clearly articulate your problem, you must first ask a series of investigative questions to build a better understanding of the prospective issue at hand and try to get at its root cause. These should focus on the **who, how, why, what,** and **when** of the issue. Working through these questions will help you to identify core issues and parse them from secondary ones.

Some questions to get you and your team started:

- What is the central problem?
    - Identify your current process or as-is state. A simple process map to outline steps (and who takes them, when) and documents/data can help you visualize the current state.
    - Identify the gaps your team is noticing that are currently not being addressed by your organization, people, and/or technologies.
    - Locate if this is a foundational problem, or just a symptom. To do this, you may need to ask "Why?" a few times to reduce the issue down to its core problem.

    - Based on the type of market research you have conducted (such desk research or an Innovator Showcase), you may have gained new insights about how GenAI's capabilities could meet the needs of your central problem.
- Who is this an issue for?
    - Map out all the key users at the different levels who are affected by this central problem. This will probably include your executives, staff/workforce, the Californians you serve, and maybe others. In what specific ways is each group of key users affected?
- How has this problem been addressed in the past?
    - Explain how you and/or others have attempted to solve this problem before, if applicable. Why were past attempts not fully successful? Past unsuccessful attempts are useful data and may influence how you write your problem statement.
- When will this problem matter?
    - Explain the timeline, how urgently a solution is needed, and why.
    - What are the risks (to your business, workforce, and/or the Californians you serve) if this problem is not solved timely?
- Why does this problem matter?
    - Explain the importance of fixing this specific problem by describing the impact it would make on the business, customers (including employees and the Californians served by your program), or other stakeholders involved.
    - It's good to think about the importance of your problem both in scale (how many people are affected? What other measurements might you have to quantify the problem?), and severity (how bad is the current issue?).
- What data do you have to support a potential solution?
    - Explain any privacy or sensitivity concerns that the data may have.
    - Explain how the data is accessed (e.g., a database) and how data accessibility may impact the potential solution.
    - Explain the overall data quality of the data available. Is it clean, structured, labeled? Think through how the data quality may impact a potential solution.

Preliminary info gathering

In order to answer these framing questions, gather as much existing research as possible to construct a more holistic picture of the issue. This can come from within your organization (e.g., business process, systems maps, org charts), or externally in the form of publications, reports, articles, etc. These can even be data points gathered from informal "interviews" with subject matter experts in the field to learn more about the topic at hand and point you in the right

direction. This is your opportunity not to reinvent the wheel by doing everything from scratch. Investigate how other people have tried to solve the problem you're facing, and use their effort to your advantage.

Data readiness considerations

If you run a proof of concept (POC) after an RFI2 procurement, innovators will only use public data when building their POC product in the sandbox. However, it's wise to think ahead to what full data you would need to build and train a robust GenAI tool if you decide to procure a full-scale product in a production environment. For a list of important questions to consider, see the Data readiness note sheet on page 4 of ODI's GenAI innovation playbook.

Equity considerations

GenAI has the potential to worsen pre-existing biases and inequities, or to create newly vulnerable or marginalized groups. If your department has an equity or external affairs lead, it's important to include them in the problem statement development process, as well as program leads. Early identification of potential equity risks in a problem statement may allow you to revise your problem statement to reduce or mitigate those risks.

Consider the following equity questions when drafting your problem statement:

- Have you talked with your equity or external affairs lead about the problem?
- What vulnerable and/or marginalized communities has your department previously worked with? What vulnerable and/or marginalized communities do your programs serve?
- Would any of those communities be affected by the proposed solution?
- Does the proposed solution amplify pre-existing risks for any groups?
- Are there alternative ways for these communities to access the service or program?

Problem statement worksheet

Using the methods above, you should begin to fill in a 1-2 page document with the following sections. You'll paste all of this content into the "Problem statement" section of an RFI2 solicitation package.

Title

Provide a simple, clear title for the problem you're trying to solve.

Overview

1-2 sentence high-level summary of the project, including your guiding question.

Ex: [Your department] is seeking to answer the problem/question of [XYZ] with a product/service that solves [describe core needs] for [your audience] by [timeframe]. This problem is important to solve because [ABC] is leading to [describe impact].

Background

1 paragraph describing the background context of your problem statement.

If relevant, consider how you or others have tried to solve this problem in the past. What worked or didn't?

Problem statement

1-2 paragraphs: Expand on the overview, and clarify the problem you are trying to solve in light of the background context. Elaborate on the ***who/how/what/when/why*** *of the problem you're trying to solve.*

Consider the following questions:

- *Who are the different users who are affected by this problem?*
- *What are they trying to accomplish?*
- *How can you make it better for the users? Without jumping to any proposed solution, what would you like to improve (e.g., faster, easier, more accurate, more efficient)?*
- *What matters most to you?*

Project goal

1 paragraph: What is your ideal ***outcome****–what would the department like to* ***achieve*** *with this solution? This is where you may speculate on what the potential solution could be for this problem, but keep it focused on outcomes and achievements. For example, a GenAI tool that summarizes vast taxpayer datasets in a fraction of the time it usually takes, so staff can better use their time to analyze data insights.*

Project objectives

If applicable, include the list of requirements/objectives that should be met by your project for success. If you include objectives, remember to keep most of them ***human-centered*** *or* ***outcome-focused****, thinking about your audience/end-users and their needs. If there are core functionalities, external data source integrations, or systems integrations that are vital to the success of a solution, you should list those items here. It's best to have as few "must-haves" as possible, so bidding innovators can pitch the most innovative ideas possible.*

In addition, the following technical and security minimum requirements are recommended for GenAI RFI2s.

The solution must also meet the following minimum requirements:

1. Vendor solution is a pre-trained, fully-functional GenAI solution.
2. Vendor will ensure the GenAI solution integrates with California Department of Technology's (CDT) managed cloud environments.
3. Vendor will host all other non-GenAI solution systems within CDT's managed cloud environments.
4. Vendor will adhere to CDT's SIMM 140 Cloud Security Guide - https://cdt.ca.gov/policy/simm/.
5. Vendor will conduct any required data pre-processing on State provided data for the purpose of tuning the GenAI solution, including scanning the data for and, as appropriate, blocking sensitive/classified data, (e.g., Personally Identifiable Information, Personal Heath Information, Federal Tax Information, etc.)
6. Vendor will monitor, report on, and as necessary, block GenAI solution prompt inputs for sensitive/classified data and any other potentially harmful prompt methods

Potential solution data sources:

- *[If you expect your solution may need to reference department-external, publicly available data sources, you can add links to a few of those sources here to help innovators begin planning their solution.]*

Sample RFI2 problem statement: Department of Finance

Title

Generating Preliminary Bill Analyses and Enrolled Bill Reports

Overview

The California Department of Finance (Finance) seeks a GenAI-enabled solution to efficiently and accurately develop legislative bill analyses utilizing existing Finance templated documents. The proposed solution should help confirm and, if necessary, correct bill information; summarize what the bill does at a high level; identify legislative redundancies within proposed legislation while populating information regarding existing laws, and the impact on the state's fiscal resources. This solution must significantly reduce the manual research workload for budget analysts, as well as integrate into Finance's existing Legislative Information System (LIS)to create a streamlined analytical process for existing budget analysts. This solution will scale insights across bill analyses to spotlight shared needs and statewide opportunities for more efficient and effective deployment of the state's budget resources.

Background

As the Governor's chief fiscal policy advisor, Finance analyzes nearly all legislation for fiscal impact and alignment with the Governor's policy priorities. This extensive workload, generally performed by Finance Budget Analysts, often competes with more urgent budget-related work. Gathering relevant background and fiscal information consumes time that is diverted from budget work. To support Finance Analysts' analyses, Legislative Affairs staff at each of the state's 150+ entities submit program-specific bill analyses to their agency, or to Finance directly upon request through separate email communications. There is no platform to standardize submissions to Finance. Finance Analysts draft their analyses in the existing LIS system by researching and utilizing these submitted department analyses, public information, previous bill analyses, and other fiscal resources.

Problem statement

Finance Budget Analysts prepare Bill Analyses and Enrolled Bill Reports (EBRs) for each bill introduced by a member of the Legislature. The main purposes of the analyses are: (1) to provide fiscal impact estimates for each bill in the Legislature, at the point in time of

appropriations hearings (Bill Analyses) or after a bill has passed both houses and is ordered to enrollment (EBR) and (2) to provide a recommended position or action to the Governor's Office. To complete these analyses, analysts must summarize the legislation, highlight current law and changes to current law, provide historical context for the purpose behind the legislation, and identify existing programs and funding from past budget bills that address the problem the bill targets. Analysts review bills and draft their analyses at the following points: bill assignment to the Appropriations Committee in the second House; and when a bill has passed both Houses and is enrolled to be sent to the Governor for consideration. This data collection and synthesis is manual and time-consuming. Unlike other state entities, Finance reviews every bill that proceeds through the legislative process which typically ranges in the low thousands.

To support Finance Analysts' analyses, state entity staff submit fiscal analyses, as well as comprehensive and templated bill analyses. Analyses include both qualitative and quantitative data, and may be shared in multiple formats (e.g., in-email text, Word documents, PDFs). This mix of submission formats creates an additional data preparation workload for Finance Analysts, making it time-intensive and challenging for Finance Analysts to identify the full scope of trends and themes across entities' analyses

Project goal

Mandatory functionality

To support staff in managing their bill analysis workload and to provide more informed and accurate analyses, Finance is seeking a GenAI-enabled solution to provide correct bill information in the relevant fields of the bill analysis form, to summarize a bill, collect fiscal information from impacted entities, to parse relevant data sets and sources for background and historical information, as well as current and emerging issues relative to a particular piece of legislation. The tool will summarize a bill and provide a high-level fiscal summary on the bills to expedite workflow for analysts. It will rapidly synthesize a significant amount of current and historical data to develop a near final analysis in Word and PDF formats, which will significantly reduce the manual workload for Finance analysts. Based on the solution-produced data synthesis, Analysts will still create the final fiscal impact estimates and EBRs.

This solution must integrate with Finance's existing LIS system, which sits in an Azure application services environment.

Optional functionality

To support Finance Analysts' analyses, state entity staff submit fiscal analyses, as well as comprehensive and templated bill analyses. Analyses include both qualitative and quantitative data, and may be shared in multiple formats (e.g., in-email text, Word documents, PDFs). This mix of submission formats creates an additional data preparation workload for Finance Analysts, making it time-intensive and challenging for Finance Analysts to identify the full scope of trends and themes across entities' analyses.

An additional, optional part of this solicitation is to develop a streamlined solution for state entities to submit their data, reduce email-based workload, and confirm receipt of their submission. For Finance Analysts, this solution would reduce manual data processing and preparation, and allow Analysts to more effectively and comprehensively identify trends and themes across entity-submitted data because the data would be centrally housed.

Project objectives

This GenAI solution will help develop and create draft Bill Analyses and Enrolled Bill Reports (EBRs) by utilizing internal files on its cloud-based servers and publicly available information to summarize and analyze proposed legislation.

Data for the solution includes public, confidential, and sensitive data, but no Personally Identifiable Information (PII) data. EBRs are confidential, and only mock versions will be provided to help vendors build a sample output.

Finance will be the sole entity to have access to the GenAI analysis tool.

If a bid includes the Optional functionality as well, the submission solution will be accessible to state staff with authentication.

The solution must:

1. Analyze a broad scope of internal and publicly available information to inform Finance bill analyses, including:
    1. Summarizing the proposed legislation as generated by LegInfo;
    2. Comparing and contrasting current law to the proposed legislative changes;
    3. Identifying previous legislation that amended the proposed code sections in the new legislation;
    4. Utilizing a standardized bill analysis template provided by Finance; and

5. Producing a draft EBR for bills that pass both Houses.

2. Integrate seamlessly with Finance's existing LIS system, which sits in an Azure application services environment.
3. Significantly reduce the time required for Finance analysts to research historical and legislative contextual information.
4. Identify redundancies between bills from different legislative sessions.
5. Provide information for each bill's current progress through the legislative cycle.

The solution must also meet the following minimum requirements:

1. Vendor solution is a pre-trained, fully-functional GenAI solution.
2. Vendor will ensure the GenAI solution integrates with California Department of Technology's (CDT) managed cloud environments.
3. Vendor will host all other non-GenAI solution systems within CDT's managed cloud environments.
4. Vendor will adhere to CDT's SIMM 140 Cloud Security Guide.
5. Vendor will conduct any required data pre-processing on State provided data for the purpose of tuning the GenAI solution, including scanning the data for and, as appropriate, blocking sensitive/classified data, (e.g., Personally Identifiable Information, Personal Heath Information, Federal Tax Information, etc.)
6. Vendor will monitor, report on, and as necessary, block GenAI solution prompt inputs for sensitive/classified data and any other potentially harmful prompt methods

Potential solution data sources:

- Department of Finance legislative analyses library
- Analysis: Principles and practices for DOF analysts
- LegInfo
- Department of Finance Ebudget Page
- Internal Department of Finance Files
- State Department Websites

Sample RFI2 problem statement: Employment Development Department

Overview

The State of California is requesting innovative ideas for a pre-trained, fully-functional generative artificial intelligence (GenAI) solution to address this Problem Statement. The Employment Development Department (EDD) seeks to address the dual challenge of enhancing recession forecasting and its ten-year industry and occupational employment projections. The EDD is seeking innovative solutions on how the use of GenAI can systematically assess and enhance, in real time, the precision of long-term employment projections and produce accurate recession forecasts to better inform state and local workforce planning and policy development. Addressing this problem is crucial, as changes in the California economy can lead to projections relying on outdated trends resulting in misinformed policy decisions, inefficient resource allocation, and missed opportunities to prepare California's workforce for future industry demands. Statistically sound methods for forecasting economic downturns are crucial because a lack of precise recession predictions limits the EDD's ability to proactively respond.

Background

For decades, the EDD has produced biennial ten-year employment projections to support California's workforce planning, policy development, and economic growth across the state and regional areas. Starting in 2025, these projections will be generated annually, increasing the need for accuracy and reliability. These forecasts are crucial for educators, job seekers, businesses, and government agencies to understand labor trends. The EDD currently lacks a systematic framework to validate accuracy of employment projections resulting in gaps at the state and local levels in understanding figures such as projected job openings by occupation, projected employment by industry, and the identification of fast-growing sectors. In parallel, the EDD plays a key role in supporting California's workforce during recessions. Accurate recession forecasting enables proactive resource allocation and guidance to impacted industries and workers. However, traditional models struggle with predictive accuracy due to complex economic variables, rapidly changing global conditions, and limited real-time data. These constraints delay the EDD's response during economic downturns and limits timely service delivery and resource availability.

Problem Statement

The EDD currently lacks a tool that aligns economic predictions with emerging employment trends, utilizes diverse datasets for refining statistical models, and supports data-driven

improvements. Additionally, it does not have a robust and adaptable forecasting system capable of using real-time data to address workforce challenges during recessions, such as resource allocation and support for impacted workers and employers. The EDD needs the ability to enhance the accuracy of industry and occupational projections while effectively anticipating economic downturns, enabling the department to respond proactively to changes in the state's economy.

Project Goal

The EDD seeks to leverage GenAI to analyze complex datasets, recognize patterns, and integrate unstructured data, thereby improving the accuracy of both employment projections and recession forecasts. The EDD's ideal outcome is to establish a GenAI solution that enhances the ten-year employment projections and provides timely, actionable insights for recession forecasting to empower the EDD to optimize resource allocation and workforce program investments. Any solution must serve as a support to the "human in the loop," enabling analysts' depth of work but not replace a person.

Project Objectives

The GenAI tool at a minimum should be able to:

1. Provide a high-accuracy recession forecasting model tailored to the EDD's needs to enable timely and precise resource allocations.
2. Provide metrics to assess the reliability of ten-year projections against actual employment data.
3. Detect patterns, identify emerging trends, as well as state and local employment shifts across industry sectors.
4. Produce a high-accuracy, enhanced projections model that utilizes real-time data (including but not limited to job postings, unemployment claims, and Worker Adjustment and Retraining Notification (WARN) notices data collected by the EDD) to supplement traditional models.
5. Provide a feedback mechanism for continuous refinement using the latest employment, economic, and demographic data.
6. Include an adaptive feedback loop for continuous improvement and mechanisms to integrate EDD analysts' expertise.
7. Provide EDD analysts with forecasting insights to promote equity, support vulnerable communities, and clarify factors driving projection variances, such as economic shifts or technological changes.
8. Provide user-friendly visualizations and reporting tools to assist EDD analysts and public and private stakeholders across the state to be able to interpret projection and recession trends in order to make informed decisions.

Innovator's solution must also meet the following minimum requirements:

1. Innovator solution is a pre-trained, fully-functional GenAI solution.
2. Innovator will ensure the GenAI solution integrates with California Department of Technology's (CDT) managed cloud environments.
3. Innovator will host all other non-GenAI solution systems within CDT's managed cloud environments.
4. Innovator will adhere to CDT's SIMM 140 Cloud Security Guide.
5. Innovator will conduct any required data pre-processing on State-provided data for the purpose of tuning the GenAI solution including scanning the data for and, as appropriate, blocking sensitive/classified data, (e.g., Personally Identifiable Information, Personal Heath Information, Federal Tax Information).
6. Innovator will monitor, report on, and as necessary, block GenAI solution prompt inputs for sensitive/classified data and any other potentially harmful prompt methods.
7. Innovator will monitor and report on GenAI solution responses for factual accuracy, coherence, and appropriateness and allow for periodic manual reviews.
8. Innovator will store all input/output data within CDT's managed cloud environments.

Recommended team roles and responsibilities for GenAI solutions

The following pages include a table of the recommended roles for a department to successfully test and oversee a GenAI solution. Although GenAI solutions are IT tools, their data inputs, training, and outputs must be reviewed by experts across program/business and IT divisions to ensure solutions perform as intended and do not produce inaccurate or biased outputs. These roles and responsibilities were successfully used for the 8 GenAI proofs of concept (POC) that California conducted in 2024-2025.

[Department] POC roles and responsibilities

POC Lead: [Name] **[Email]** **[Phone]**

	Primary team						
	POC Lead	**Product**	**Project**	**Labor/ Workforce**	**Technical and Security**	**Data**	**User experience research (UXR)**
Assigned	**Name** Title Email Phone	**Name** Title Email Phone	**Name** Title Email Phone	**Name** Title Email Phone	**Name** Title Email Phone	**Name** Title Email Phone	**Name** Title Email Phone
Role purpose	**The single point of contact for the GenAI POC.** The POC Lead oversees all other POC role areas to keep a 360° view of POC progress, mitigates known risks, and flags	Each GenAI tool being considered must be evaluated holistically, including a thorough assessment of the potential tool's fit for a department's program, operations, staff,	POCs are run on an approximately 6-month time frame. The Project Manager works with the innovator and internal teams to ensure timelines and milestones are	Independent of cost and technical concerns, a POC must support and deliver value to the State staff who use it. The State staff who will be end-users of the POC tool must be involved early and	A learning outcome of a POC is to anticipate the production needs of a procured GenAI tool. The Technical and Security Lead will support GenAI production planning, including	GenAI tools depend on high-quality data that can be leveraged for tuning and analysis. A subject matter expert (whether from Program or IT) must be familiar with the data needed for	GenAI tools have the potential to improve the speed, accuracy, and scale of what state workers produce, freeing them from more manual tasks so more time can be spent on analytical

	concerns for leadership as they arise so they can be addressed (including potential policy updates as necessary). **For overall POC success, ODI recommends that the Product Lead also serve as the POC Lead.**	and the Californians they serve. The Product lead is responsible for the human-centered focus of the POC.	met.	often to develop test cases of the tool, provide feedback on its ease and accuracy of use, and evaluate its success as a product. In coordination with the department's Labor Relations Officer, the Labor/Workforce Lead ensures that staff's opinions, use cases, and evaluations drive the product assessment of the POC. They also ensure appropriate notification to CalHR as needed for potential meet and confer triggers.	what is required to set up and optimize a technical infrastructure for a production environment.	the specific use case, with insight into the data collection, refinement, and maintenance details of each data source. This person should work with ODI and CDT as necessary to assess the readiness of potential data sources, and co-develop risk mitigation plans, for both a POC/sandbox and future-state production tool. If data mitigation (including imputation, enrichment, and/or redaction) is needed, this person will assist in those processes.	or creative work. GenAI tools must directly benefit the state staff who use them. Conducting user experience research (UXR) with the state staff who would use the POC tool is the means to determine whether the tool meets their needs and improves their daily work. UXR also creates a backlog of future-state builds and enhancements is that can drive the ongoing product roadmap if the POC tool is procured. UXR may be conducted through interviews, surveys, or a combination.
Depart-mental role	Senior leader (historically, has been the Product	Senior business/program leader in the	Assigned Project Manager	Mid-level manager in the division where staff will use	Senior leader(s) in the IT division who can provide	A mid-level or senior leader with awareness of the	A mid-level or senior leader who has taken ODI's

	Lead as well: a senior business/program leader in the division where the POC use case is being tested)	division where the POC use case is being tested		the POC tool, who can represent staff views and knows which staff to assign for POC evaluation	expertise to the POC Lead	department's proprietary, program, and public data– including its known issues or shortcomings	UXR training. This person will partner with the Labor/Workforce Lead to identify staff for UXR interviews
Reports to (for POC purposes)	Your department's Executive office	POC lead	POC Lead	POC Lead	POC Lead	POC Lead	POC Lead
Responsi-bilities	Serve as the single point of contact for their internal team. Ensure execution of all milestones on POC product management roadmap.	Liaise with Department and Agency Executives to evaluate POC's success from a human-centered and customer perspective.	Support execution of all milestones on POC product management roadmap.	Lead workforce/labor testing of the POC tool, including test question development. Support execution of all milestones on POC product management roadmap.	Partner with CDT to assess data security, PHI/PII, etc. Monitor to ensure sandbox utilization is as expected. Closely monitor inbound/outboun d of data and outputs. Ongoing security monitoring.	Partner with ODI and CDT to assess data readiness, governance, and completeness. This assessment will occur before the POC sandbox launch. Throughout the POC phase, closely monitor inbound/outboun d of data inputs and outputs, as well as performance monitoring for the tool.	Partner with internal leaders to identify UXR interviewees, schedule, and hold UXR interviews.

	Wraparound support team			
	Legislative Affairs	**External Affairs**	**Legal, Privacy, and Procurement**	**Communications**
Assigned	**Name** Title Email Phone	**Name** Title Email Phone	Legal: **Name** Title Email Phone Privacy: **Name** Title Email Phone Procurement: **Name** Title Email Phone	**Name** Title Email Phone (This individual may be the same as the External Affairs Lead, depending on your organization's structure.)
Departmental role	Senior leader in the Legislation division who can provide expertise to the POC Lead, and interface with your agency's Deputy Secretary of Legislation	Senior leader in the External Affairs division who can provide expertise to the POC Lead, and interface with your agency's External Affairs team	Senior leaders in the Legal, Procurement, and Policy or IT divisions (as applicable) who can provide expertise to the POC Lead	Senior leader in the Communications division who can interface with your agency's Comms team
Reports to (for POC purposes)	POC Lead	POC Lead	POC Lead	POC Lead

How to run a pre-mortem meeting

At the end of a project or sprint, it's a good idea to conduct a retrospective to look back on what did (and didn't) go well so you can improve. But a "retro" is just that—looking backward at what has already been done. It's reactive.

But at the start of any important or complex project, such as a GenAI procurement, it's wise to be proactive. This means identifying actual and potential risks or threats so they can be avoided if possible. This proactive approach is called a pre-mortem. **Holding a pre-mortem increases your chances of project success.**

This document teaches you how to run a pre-mortem meeting.

Participants

In order to reap the greatest benefits from a pre-mortem, you should include all members of your POC product team (see the Recommended team roles and responsibilities for GenAI solutions tables beginning on page 22 of the ODI GenAI innovation playbook). You should also include leadership and executives, including the directorate and even your agency if they're available.

At a minimum, you'll want staff from the program area (including subject matter experts in the data that will be needed for the GenAI solution), IT, security, Legal, Communications/External Affairs, procurement, project management/PMO, and HR. Including a broad mix of business areas will help you capture all risks and identify mitigations.

Logistics: location, supplies, length, roles, and documentation

The best pre-mortems happen in-person, using sticky notes. You should have enough sticky notes for each participant to write 5-10 notes. Allowing participants to write their answers and place them around a room encourages honesty and full participation.

If an in-person meeting isn't possible, you can use a digital sticky note alternative such as the Whiteboard feature in MS Teams, Mural, Miro, Figma, or similar software. If it's possible to let people add virtual sticky notes anonymously, you should do that.

You should dedicate at least 30 minutes to a pre-mortem. Large groups, or particularly complex projects, may need 60 minutes to allow for a more robust discussion.

To run a pre-mortem, you'll need at least one moderator and one note taker.

Regardless of whether your meeting is in-person or virtual, you'll want to share a slide with instructions. *(Recommended language is included at the end of this document.)*

You'll also need to create sorting space for the (paper or virtual) stickies. For an in-person meeting, this can be written on a whiteboard or on extra-large wall sticky notes. For each category (Tigers, Paper Tigers, Elephants), you'll want to create a sorting "T" like this:

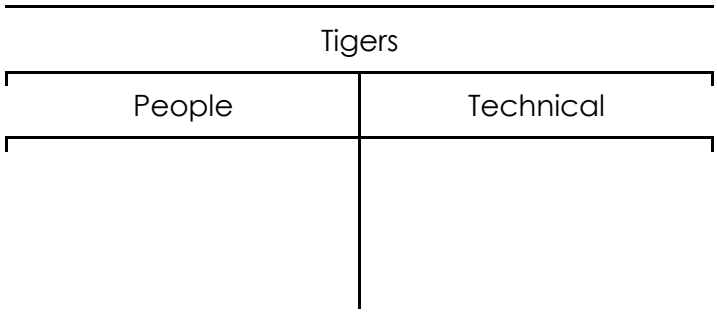

Definitions

- **Tigers**: a threat you are afraid of
- **Paper tigers**: threats that others might worry about, but you are not afraid of (due to scope, because it's within your program area to mitigate, or for some other reason)
- **Elephant**: a failure no one wants to say out loud–but folks need to be aware of

Pre-mortem run-of-show

1. **Assign roles:** To run a pre-mortem, you'll need at least one moderator, and one note taker.
2. **Welcome:** Identify all the different business areas present, and allow participants to introduce themselves, their title, and role. Especially in larger departments, folks may not know everyone present or how they relate to the project.

3. **Set the stage:** Give context for what you're about to do. By proactively thinking about how a project could fail, you're identifying potential risks and their mitigations. You're raising awareness of "red flags" for the entire group to be aware of and collaborate on. You're planning for success. *(Suggested slide deck content is at the bottom of this document.)*
4. **Explain the terms:** Using the definitions above, explain what Tigers, Paper Tigers, and Elephants are. Explain that you'll be sorting stickies by "People failures" and "Technical failures."
5. **Logistics:** If you're in-person, pass out sticky notes and pens. If you're virtual, explain how to use the whiteboard software.
6. **Share the ground rules:**
    a. Be honest. The team can't mitigate risks that aren't shared, and many risks will take a cross-functional team to avoid or mitigate effectively.
    b. You don't have to be an expert to identify a risk. "Gut senses" are ok in a pre-mortem. For example, an IT team member could identify a Tiger about how this project will be communicated externally. Or Comms team member could identify an Elephant about business data, and so on.
    c. Be generous. The nature of a pre-mortem is that it will raise awareness of risks and weaknesses in your department. Some of these risks will "belong" to one division or team, but **identifying** a risk is not the same as **blaming** the team responsible for the current state.
7. **Brainstorming portion:** Begin the brainstorming. Use a timer. For a 30-minute meeting, you should give 7-10 minutes for brainstorming. Encourage participants to write at least 2 each of Tigers, Paper Tigers, and Elephants. (Although most people will write much more than 2!)
8. Sorting: Participants will put
9. **Reading portion:** Give participants a few minutes to read all the sticky notes.
10. **Paper Tiger resorting:** If a participant sees a Tiger or an Elephant that they know is (1) able to be fully avoided or mitigated and (2) lies within their responsibility, they can move a sticky to the "Paper Tiger" section.
11. **Discussion portion:** Once everyone has had the opportunity to read all notes, open the meeting up for discussion. As a group, use the remaining 20 minutes to walk through the common themes and discuss what emerged. Some potential prompting questions are:
    a. What stood out to you?
    b. Were you surprised by any stickies?
    c. Did you notice any themes?
    d. Which are the most important?
    e. Why are these coming up now?

f. How might we address these risks?
g. Did anyone move a sticky to the "Paper Tiger" section? Share more about that.
h. Did the same issue come up as both a Tiger **and** a Paper Tiger? If so, have a discussion about why some participants believe the threat to be a Paper Tiger. Are there mitigations already in place?
i. *(Choose an Elephant sticky–perhaps the most important, complex, or volatile one.)* Let's talk about this Elephant.

12. **Take notes:** Have your notetaker keep track of them and the discussion. Use these discussion notes as a foundational component of writing your team's problem statement.
13. **Wrap-up:** Thank everyone for attending. Explain that the data gathered today will be typed up and shared with all participants. *(A suggested note format is at the end of this document.)*

Post-meeting notes

After the meeting, the notetaker should capture all sticky notes and where they were sorted into a document. You should watermark this document as "Draft–deliberative" or similar. It is a best practice to share this document with all pre-mortem attendees.

Below is a recommended table format for memorializing all sticky notes and their final placement:

	People	Technical
Tigers	●	●
Paper Tigers	●	●
Elephants	●	●

Suggested slide deck content

If you would like to have an instructional slide showing during your meeting, you can use the language below:

Let's imagine this project does not succeed on the technical side or people side.

Why would that happen? What did–or didn't–happen to lead to a proof of concept failure?

1. **Brainstorm**

Each person should brainstorm:

- **Tigers**: a threat you are afraid of (at least 2)
- **Paper tigers**: threats that others might worry about, but you are not afraid of (due to scope, because it's within your program area to mitigate, or for some other reason) (at least 2)
- **Elephant**: a failure no one wants to say out loud–but folks need to be aware of (at least 1)

2. Sort

As you place your tigers/paper tigers/elephants up, identify if they are **People failures** or **Technical failures** and place them in the right column