Step 1: Pinpoint high-priority operational challenges

↳ Consider operational processes that cause friction, then use the flow diagram to identify your high-priority challenges

STEP 1: OPERATIONAL CHALLENGE PRIORITIZATION

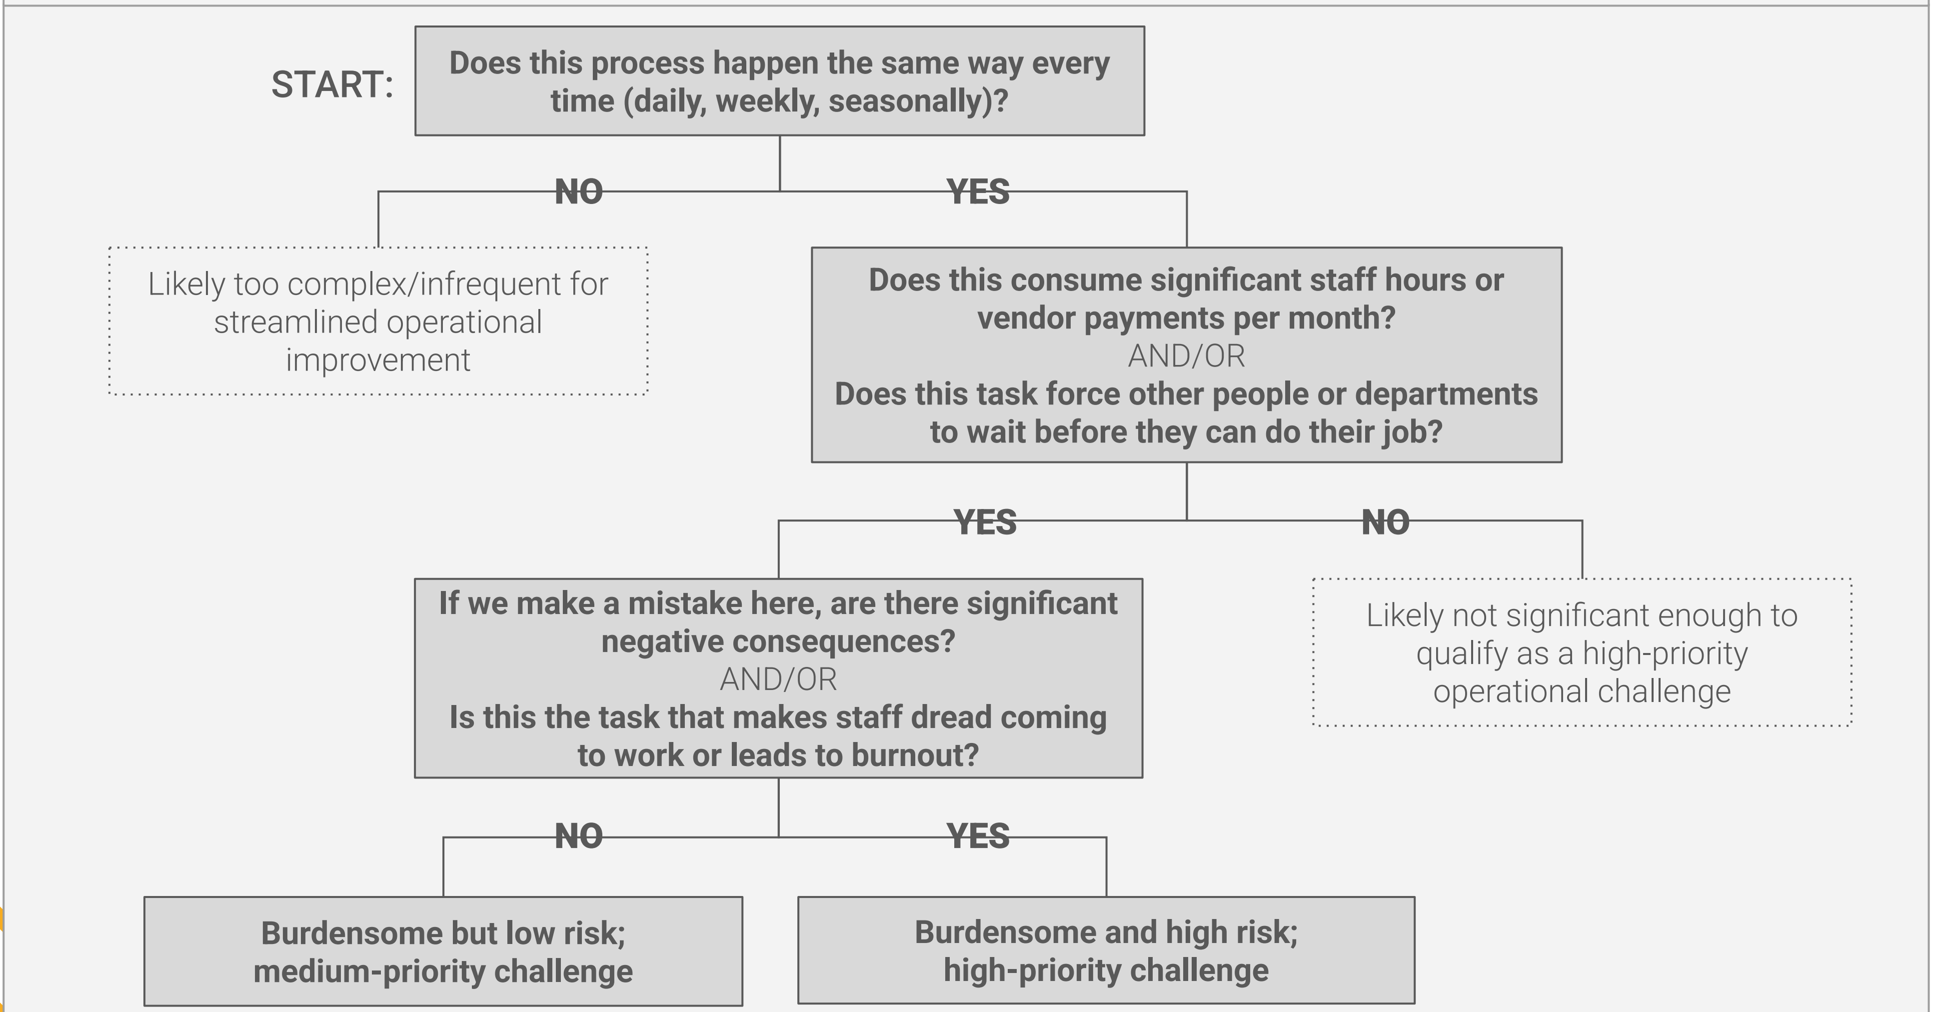

Step 2: Assess whether data science & AI can help address your high-priority challenges

↳ Consider the workflow of your high-priority challenges, then evaluate each challenge against the matrix below, sidelining those that are not viable candidates for data science or AI augmentation

<table>
<tr><th colspan="4">Step 2: Data Science and AI Viable Use Cases</th></tr>
<tr><th></th><th>Strong Case for Augmentation</th><th>Some Potential for Augmentation</th><th>Human-Critical Tasks</th></tr>
<tr><td>Generative AI</td><td>• Info search, gathering, & synthesis
• Document & report drafting
• Communication & correspondence
• Interpretation & translation
• Training & instruction preparation
• Help desk chatbot
• Code assistant</td><td>• Decision-making & prioritization
• Strategic planning
• Hiring management
• Ethical/legal review
• Database management
• Brainstorming
• Quantitative analysis</td><td rowspan="2">• In-person relationship building
• Complex negotiation
• Ethical leadership & accountability
• Institutional representation
• High-stakes/high-risk decisions</td></tr>
<tr><td>Data Science</td><td>• Forecasting & demand planning
• Early warning signals and error detection
• Logistics & resource optimization
• Quantitative analysis & inference
• Pattern recognition & categorization</td><td>• Data cleaning & preparation
• Scenario simulation
• Causal inference
• Defining success metrics</td></tr>
</table>

Step 3: Navigate the complexity and risk of solutions driven by data science & AI

Use the framework below to assess complexity and risk of your potential data science & AI solution(s). Consider the questions (Slide 3) and their associated details (Slides 4-6) to gain insights into the complexity and scope of your initiative. For each dimension, examples with varying implications are provided.

STEP 3: DIMENSIONS OF DATA SCIENCE & AI IMPLEMENTATION		
1. Strategic Intent	**1. Autonomy**	**1.1.1** Will the solution provide recommendations for human decision-making, or will it take actions autonomously? **1.1.2** Are there tasks that must remain under human control for regulatory, ethical, or risk management reasons?
	2. Mode	**1.2.1** What types of input data will the solution need to process (numeric data, text, images, audio, video)?
	3. Function	**1.3.1** Is the primary goal to streamline a task or render a process more efficient? **1.3.2** Is the primary goal to analyze what happened/is happening, or to forecast or optimize what will happen?
	4. Explainability	**1.4.1** Will stakeholders need to understand how the solution reaches its conclusions?
2. Technical Implementation	**1. Architecture**	**2.1.1** Are there existing out-of-the-box solutions that meet your needs, or does the solution require custom integration? **2.1.2** Do you have the expertise within the team to architect, cost-out, and build a solution in-house, or manage a vendor?
	2. Data Context	**2.2.1** Will off-the-shelf models trained on general data suffice, or does the solution require specific domain knowledge? **2.2.2** Will proprietary or personally identifiable data (PII) information need to be processed as part of solution?
	3. Data Accessibility	**2.3.1** Where does the required input data currently reside, and how easily can the data be accessed?
	4. Data Format	**2.4.1** Will your solution require input data that is consistently formatted and complete, and if so, is significant cleaning and preprocessing required?
	5. Maintenance	**2.5.1** Will this solution's performance degrade overtime without regular monitoring and maintenance?
3. Scope	**1. Scale**	**3.1.1** Will the solution address a single-user specific function, or will it have cross-organizational implications? **3.1.2** Will the solution run a data science or AI model with high and regular frequency, or only as-needed?
	2. Rollout	**3.2.1** Who will be the primary users, and what is their technical sophistication? **3.2.2** Does the organization currently have in-house technical experts that can assist with user questions and/or engage with/provide feedback to a vendor?

STEP 3: ADDITIONAL DETAIL

1. STRATEGIC INTENT

1. Autonomy	**1.1.1** Will the solution provide recommendations for human decision-making, or will it take actions autonomously? **1.1.2** Are there tasks that must remain under human control for regulatory, ethical, or risk management reasons? As autonomy increases, the human role shifts from "doing" to "auditing." Tasks delegated to AI could result in greater efficiency gains than AI serving as an assistant; but, human review and monitoring are critical for ensuring accuracy and compliance. Large scale, complex workflows pose real technical challenges in accuracy and monitoring, especially in cases of outliers or failures. **Consider:** a tool suggesting interview questions for a new teacher candidate based on the job description, **versus** a system that automatically texts substitute teachers to fill vacancies in real-time without human intervention. Any tasks that involve high-stakes personnel decisions, student discipline and due process, resource equity, or crisis management, as well as those that are in direct violation of data privacy and compliance regulations, should remain under human control.
2. Mode	**1.2.1** What types of input data will the solution need to process (numeric data, text, images, audio, video)? Text reading and writing is cheap and fast, while more specialized or multimodal uses could require the use of more expensive, slower models and associated tools; video and audio models are typically most expensive per task. **Consider:** a tool that specializes in html coding for website management, **versus** a system that analyzes school security camera feeds for door-propping or safety hazards. Any task requiring a factually correct answer (a statistical analysis, an executed script of code) requires human oversight and frequent spot checks.
3. Function	**1.3.1** Is the primary goal to streamline a task or render a process more efficient? **1.3.2** Is the primary goal to analyze what happened/is happening, or to forecast or optimize what will happen? If the goal is to streamline a task or render a process more efficient, success metrics will vary based on the objective of the process, but should be tied to real-world impacts (hours saved, dollars spent). If the goal is to analyze, forecast, or optimize, the implementation requires mechanisms to measure accuracy over time and communicate uncertainty to users. **Consider:** a function that categorizes emails sent to the superintendent's office and automatically routes them to the correct team, **versus** a function that predicts next year's kindergarten enrollment based on local demographic data.
4. Explainability	**1.4.1** Will stakeholders need to understand how the solution reaches its conclusions? Transparency can be a crucial element of many initiatives, both for stakeholder trust and for explainability. Higher-stakes decisions should prioritize explainability over pure accuracy or function to avoid bias and maintain community trust. **Consider:** a tool that automatically translates press releases into multiple languages, **versus** a rule-based system flagging students for intervention based on clear and explainable metrics and thresholds.

STEP 3: ADDITIONAL DETAIL
2. TECHNICAL IMPLEMENTATION

1. Architecture	**2.1.1** Are there existing out-of-the-box solutions that meet your needs, or does the solution require custom integration? **2.1.2** Do you have the expertise within the team to architect, cost-out, and build a solution in-house, or manage a vendor? Standalone tools or external tool connections present lower barriers to entry in terms of technical requirements and complexity, but are less customizable, and can introduce greater risk of vendor lock-in. Further, users still need to be trained on their appropriate use, especially open-ended GenAI chat-based tools. Deep integrations are more technical, requiring API and/or open-source builds and (on-going) engineering resources. **Consider:** leveraging a third-party tool (like Zapier) to send an AI-generated summary from a Google form to a Slack channel, **versus** an AI layer embedded directly into the Student Information System that triggers actions in real-time based on database changes.
2. Data Context	**2.2.1** Will off-the-shelf models trained on general data suffice, or does the solution require specific domain knowledge? **2.2.2** Will proprietary or personally identifiable data (PII) information need to be processed as part of solution? Familiar GenAI models operate in a generic sense, without specific training or expertise in specialized domains. Some out-of-the-box tools and apps tailor models specifically for distinct domains (like legalese). Models can also be fine-tuned in-house on proprietary data for this purpose. **Consider:** summarizing public board meeting minutes, **versus** interpreting and composing legal documents like vendor contracts or leases. Solutions that requires access to proprietary data or PII could require fully internal infrastructure (using open source resources) or, at minimum, external tools & APIs with rigorous compliance with data privacy standards, including role-based access controls.
3. Data Accessibility	**2.3.1** Where does the required input data currently reside, and how easily can the data be accessed? Data silos are a significant challenge; siloed data likely require a custom in-house build to connect to AI technologies or to be fed into a data science model. Real-time data stream requirements also pose a technical integration challenge, at minimum requiring API maintenance and monitoring. **Consider:** an AI tool that must interact with a specific HR-based application to respond to payroll inquiries, **versus** a system that integrates live GPS feeds from buses with an AI model to text school leadership updates on potential late arrivals.
4. Data Format	**2.4.1** Will your solution require input data that is consistently formatted and complete, and if so, is significant cleaning and preprocessing required? Data format is more relevant for solutions that require strict input structure (forecasting, clustering) and less so for GenAI solutions, which are more flexible. Complex data preparation and cleaning can comprise more time and effort than the solution itself, so mapping current state is critical. **Consider:** a budgetary shortfall prediction model based on utility costs and staff overtime metrics, **versus** a chatbot that can respond to questions.
5. Maintenance	**2.5.1** Will this solution's performance degrade overtime without regular monitoring and maintenance? Unlike static software, data science and AI models "drift" as the real world changes (e.g., traffic patterns change, affecting bus routes). Districts must account for the long-term cost of monitoring the accuracy of <u>any</u> solution, but particularly those in-house builds that rely on dynamic, context-specific data. **Consider:** a tool that extracts text from scanned forms or PDFs, **versus** a system that identifies students at risk of chronic absenteeism or course failure based on attendance patterns, grades, and engagement metrics.

STEP 3: ADDITIONAL DETAIL 3. SCOPE	
1. Scale	**3.1.1** Will the solution address a single-user specific function, or will it have cross-organizational implications? **3.1.2** Will the solution run a data science or AI model with high and regular frequency, or only as-needed? Individual use cases present “soft” time savings, which can be difficult to measure. That said, they are often lower risk and can be adjusted based on focused user feedback. Widespread tool usage can be easier to measure impact, but often have high up-front costs and require cross-functional alignment and executive sponsorship. Complex system reconfigurations require a dedicated Project Manager and change management strategy. **Consider:** an AI assistant that helps a staff member compose emails, **versus** a system that continually monitors HVAC, electrical, and plumbing systems in real-time across all district buildings, predicting likely equipment failures and and automatically scheduling preventive maintenance. Further, particularly for API calls to external data science & AI models, the frequency/volume of the calls is directly related to cost. The more often the model is run, the higher the cost. This is especially relevant for automated high frequency requests, such as real-time chatbots fielding thousands of daily queries.
2. Rollout	**3.2.1** Who will be the primary users, and what is their technical sophistication? **3.2.2** Does the organization currently have in-house technical experts that can assist with user questions and/or engage with/provide feedback to a vendor? Training in appropriate and effective use of AI is a significant consideration, even for open-ended chat-based tools. Individuals and teams directly leveraging AI tools must be trained on their appropriate use and limitations; staff should be empowered to override AI output and recommendations when concerns for accuracy or hallucinations arise; and each initiative should have a known technical “owner” or point-person with whom concerns or questions can be raised. **Consider:** a single school staff member leveraging an AI tool to draft the weekly "Friday Folder" newsletter from bullet points, **versus** equipping hundreds of bus drivers with an AI-embedded tool that suggests route adjustments based on morning traffic. Data science & AI implementations must be designed to accommodate varying skill levels among those tasked with using them, with intuitive user interfaces. Monitoring adoption metrics are critical here, as an unused tool or system is not worth the investment.

Step 4: Brainstorm tangible and measurable indicators of success

Fill in the value proposition statement to begin weighing considerations and complexities against impact value

If we could reduce the

. . . time spent on . . .
. . . money spent on . . .
. . . errors produced within . . .

this task by ______________ %, then our team would be able to

. . . redirect our focus toward ______________.
. . . reduce our monthly expenditures by ______________.
. . . improve our staff retention rate by ______________.