

bigspark's AI Advisory - Context and delivery
guide - v0.1

Goals, process and accelerators

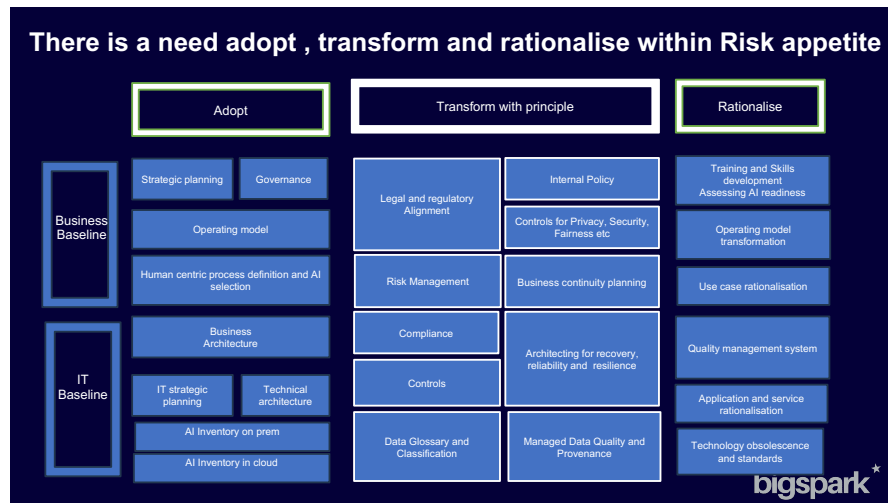
Table of Contents

<i>bigspark's AI Advisory - Context and delivery guide - v0.1</i>	1
Goals, process and accelerators	1
Goals - v1.0	3
What the client wants	3
Why create automated accelerators?.....	3
AI Advisory goals and target clients	5
Process (Step 1) – Intention workshop	6
Accelerators (Step 1)	16
Audiences	16
APPENDIX –	20
Bigspark's Knowledge Base Accelerator specification	20
User Personas	22
1. Data Steward (Maria, Kate, bigspark best practice leaders and founders):	22
2. AI Advisor:.....	22
3. Administrator:	22

Goals – v1.0

What the client wants

Clients seek to achieve one or more aspects of AI adoption, literacy and transition as characterised below. Success metrics will be aligned with the client's view of scope and priority.



Why create automated accelerators?

The tool suite we are developing is designed to assist AI advisors in the process of advising clients on AI adoption and literacy through ongoing AI lifecycle management. Use of a knowledge base and a consistent tool set makes it easier to scale bigspark's AI advisory service by creating a 'formulaic' lego block approach.

The various stages of interaction anticipated when engaging a client are shown below. Please note that it is not always necessary to start with an intention workshop but given its pivotal role we are starting collateral development with this in mind.

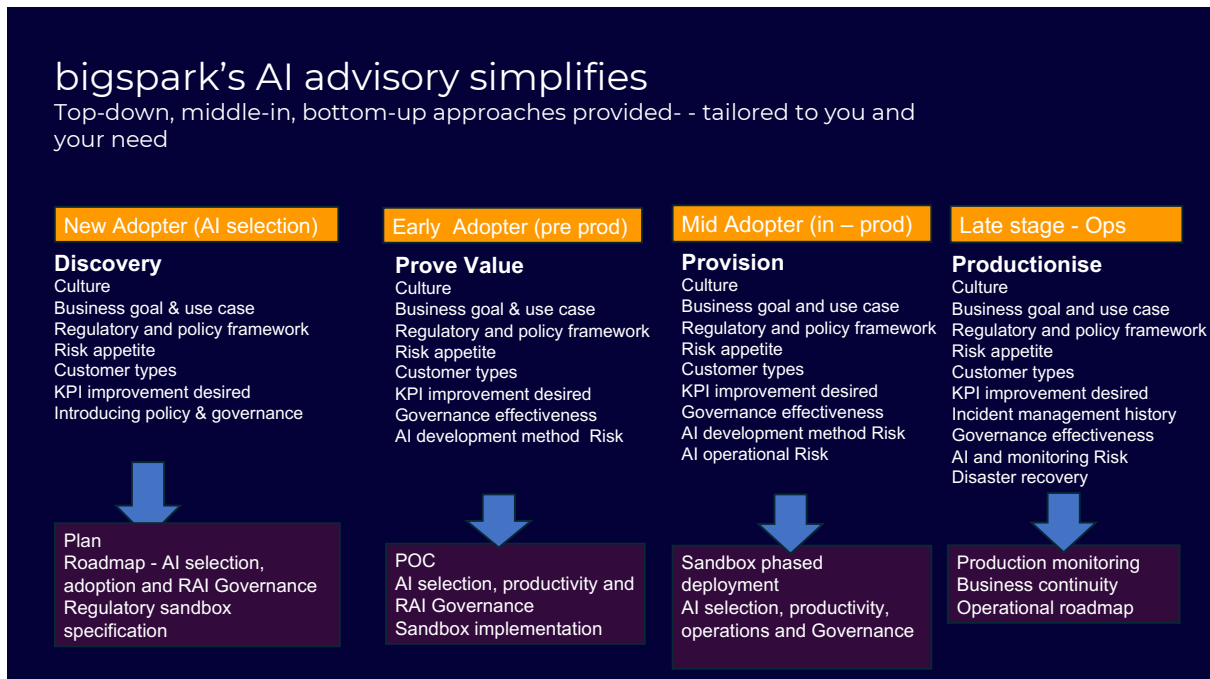


As indicated it is our goal to curate, map and advise a client to plan AI adoption or improved AI usage. Our Responsible AI work is always traceable to regulatory and legal constraint within jurisdiction and industry and the use of our 2 + 10 principles are used as a 'translation lens' to make it possible to trace from law/regulation/policy, through principle, to use case and plans.

The goal of step 1 (AI advisory bootstrap) as shown above is to identify the operating context of the client. Know their perceived risk and reward intention and then help them manage their Governance, Operating model and ongoing implementation in a way that aligns with this intention for risk appetite and reward. Metrics are critical. Please note that it is beyond the scope of this document to discuss automation requirements of step 2 and step 3 as shown in the above graphic. In synopsis Step 2 will focus on advising clients on how to leverage sandboxes and ongoing governance techniques to validate ongoing AI conformity and step 3 will enable Third parties to integrate. Eg. Audit review, Information sharing etc.

AI Advisory goals and target clients

AI Advisory on adoption, management and literacy can be provided to clients in any stage of the AI adoption lifecycle as indicated below. AI adoption is a horizontal endeavour and not linked to any specific industry. It is a fact of course that any one client might employ multiple AI adoption and management strategies at once



The goal of bigspark's AI advisory is to provide value to clients throughout. Evidently bigspark benefits by supplying engineering services that address step 2 and 3 considerations. It is the job of the AI Advisory to build a relationship with the client and pull through ongoing Consulting and Engineering services.

Process (Step 1) – Intention workshop

The first step on engaging the bigspark Advisory is focused on the delivery of an intention workshop. Here's what preparation and steps are employed to perform an intention workshop. Items in red are candidates for the public Microsoft RAI companion app

- 1. Build and enrich a bigspark knowledge base: Advisors constantly curate and gather public domain information mapped to rules and questions** (View 1). Requirements are gleaned from LRPS (Law, regulation policy and specification inputs or/and bigspark and authoritative technology practitioners). View 1 also includes best practices and guidance rules from bigspark prior engagements. Finally View 1 includes items related to bigspark's AI advisory methodology such as the definition of quadrants. Management and maintenance of the knowledge base is ongoing (*Note we need more input from technical bigsparkers- we should agree how to pursue*)
- 2. Client introduction: context setting and preliminary information capture:** When meeting a client Advisors need to capture client specific information. The client will have specific use cases, will have specific KPIs, client policies, priorities etc. Bigspark advisors must query the client to gather this input. **At the start of any client engagement – and before pricing and intention workshop scope can be agreed – an NDA must be signed and the client must be asked to articulate (i) their top three Risk and their top three reward perceptions of AI adoption. The client must also be asked to provide (ii) on set of prioritised AI adoption/usage use cases (iii) a statement of industry scope and (iv) a statement of jurisdictions to be considered. This is to be done through the public offering on the Microsoft Marketplace**
- 3. Intention workshop pricing and paper signing:**

Based on the input provided in Client introduction, a price for the advisory is calculated. A limit of 10 use cases is to be imposed by client but the exercise can be repeated again and again. We also encourage a client to pick one industry and one jurisdiction to start. In order to conduct pricing estimations The bigspark advisors map use cases onto stereotypes and place them in the bigspark AI management quadrants to form a hypothetical assertion of the 'spread' of use case and AI management needs. This is a bigspark internal only activity. The greater the spread the higher the likely cost. If there are no pre-existing stereotypes then the

higher the cost etc. Here is a view of the AI risk assessment quadrants being adopted.

				Secure				
		Quadrant A			Quadrant B			
				3				
				2				
Open Data	3	2	1	1	2	3	Highly classified Data	
				1				
				2				
				3				
		Quadrant C			Quadrant D			
				Public				

4. Intention workshop preparation:

When a client has signed and committed to pay for the intention workshop the team prepare. The use cases mappings to existing 'best AI practice' **stereotypes and/or new stereotypes are created**. The hypothesis can now be shared with the client. Stereotypes are typical best practices yielded by the use of AI in a quadrant (eg if on prem network isolation is required for an LLM (typical in the Quadrant B) then a ChatGPT endpoint is not a viable option. To this end Chat GPT cannot generally be used in Quadrant B but Meta llama or BERT can be. To this end bigspark has, from the outset an assertion of a high level checklist of best LLMs (and SLMs, and RRTs) fit for purpose by quadrant and to this end bigspark can compile lists of AI governance and management methods likely to be required.

Each stereotypes has known attributes and known 'best practice LLM affiliates (Note Dimple is working on constructing this). As an example here is a stereotype summary of a view focused on guiding clients on the use and adoption of AI embedded in vendor tools – Notice there's not yet a use case so we don't know what Vendors are being used

precisely but that said there are already a set of things that we're going to have to address and advise on.

The AI You Didn't Know You Had

Procurement

Use Case 1:

Stereotype:

Third-party tool vendor risk management

Jurisdictions:

England & Wales

Focus:

- Address specific rogue unacceptable use cases
- Add any contractual terms needed to govern AI to existing MSAs, SOWs, etc
- Adjust procurement policy
- Introduce minimum sufficient 'safeguard' Policy to manage detected Sizewell C AI usage.
- Improve AI literacy.
- Reduce existing liability
- Ensure removal of prohibited use cases detected using EU AI Act terminology. bigspark validates against EU definitions of AI risk in all jurisdictions.
- Introduce inventory, test, audit and active Governance of detected High risk AI



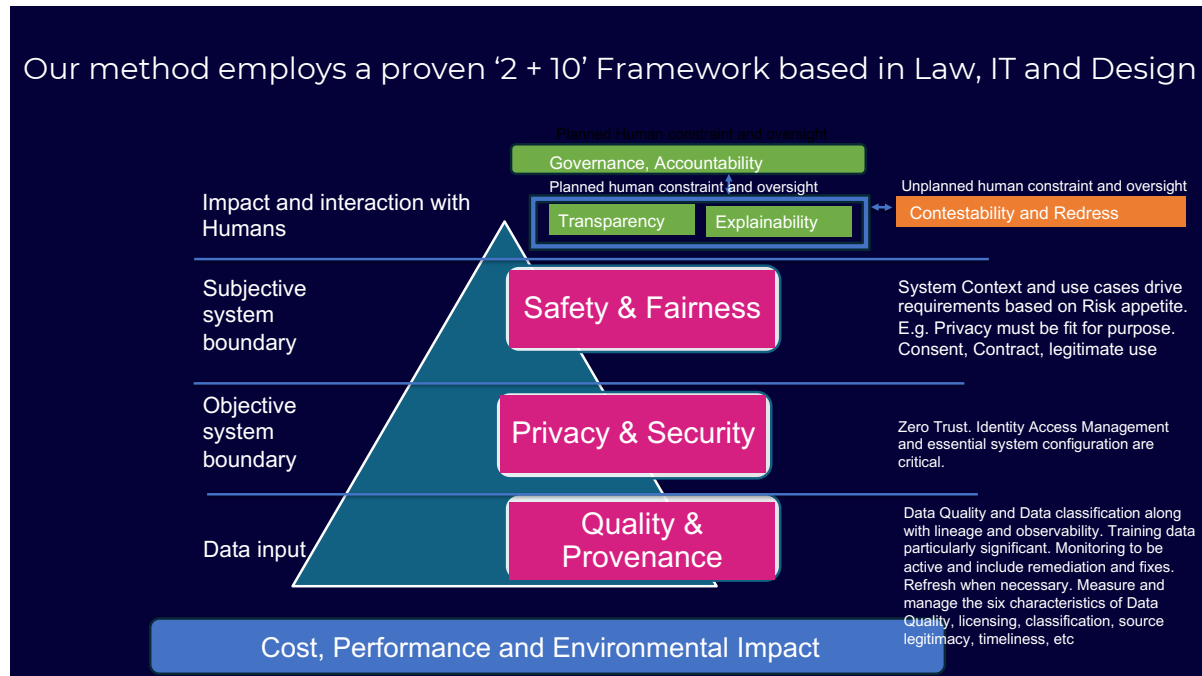
Assumed Data: (Tick all that apply)

- Public Data
- Internal Data
- Confidential Data
- Restricted Data
- Private Data
- Critical Data
- Regulatory Data

User Groups: (Tick all that apply)

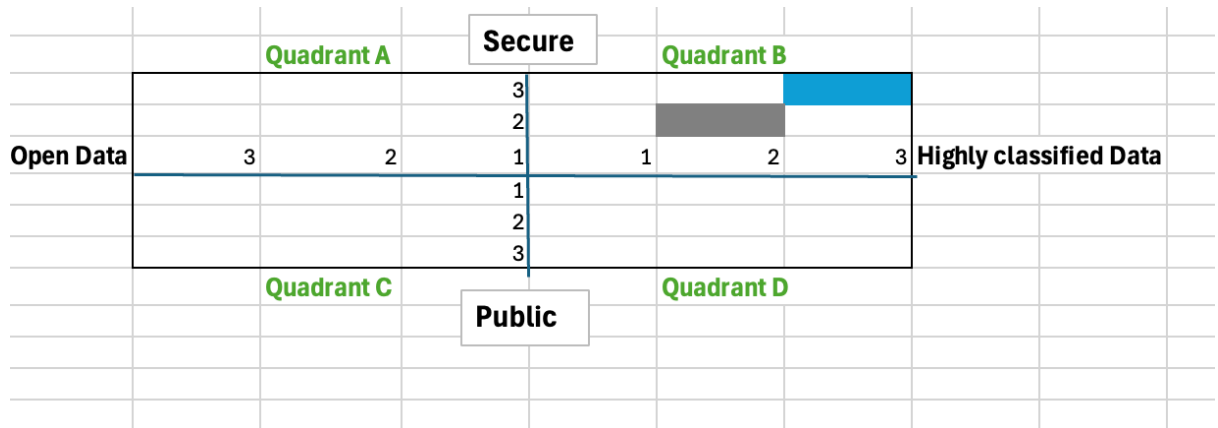
- Secure.users.Administrators
- Authorised users.Administrators
- Authorised users.Legal
- Authorised users.Auditors
- Authorised users.Employees
- Authorised users.Suppliers
- Authorised users.Management
- Authorised users.Temporary assignees
- Guest.Businesses
- Guest.Legal
- Guest.Individuals
- Other

5. **Creating a client questionnaire:** bigspark draws upon its knowledge base by default. If there are any regulations or laws relevant to the client use case then we use the '2 + 10 principle lens' to select the rules and questions from the curated list of LRPS approved by the client's legal counsel by default from the bigspark knowledge based. If the client asks for new laws to be included then we prepare the rule supplement for this new law. To this end the fundamental and essential questions about responsible AI principle requirements, (e Privacy, Security and Fairness etc) are derived from any law that applies. See graphic below as reminder of the principles



In order to ask more nuanced questions, tailored to the context of the client, we map use cases (through their association with the stereotype) to the Bigspark's risk assessment quadrant table. As an example might be that the client seeks to use AI to find contractual clauses of concern regarding the legal requirements of building nuclear reactors. This information is to be considered sensitive and cannot and must not leak or be shared in the cloud. Equally well, this same stereotype of law and guidance might also support a Human Rights Law Firm. The use cases are different but the stereotype – Legal confidential data management - is largely constant. For the sake of illustration we might agree that while the stereotype is similar we do want a slightly different profile applied to this

one stereotype of legal confidential data management. One profile for B2,2 (law firm) and one profile for B3.3 (nuclear legal liability).



While the **stereotype and profile** are ‘View 1 assets – aka reusable , the **mapping of use case to stereotype** is a view 2 asset as the mapping of the stereotype to the risk matrix is tailored for the client.

To further illustrate what is said above the following bullets attempt to exemplify what we would do using the real example of Sizewell C.

- Sizewell C have chosen to focus on managing use and choice of AI by Microsoft and Adobe tool vendors. By contrast, a law firm might be focused on whether or not to use CaseText or IronClad Legal Document management tools. This shows that while we will always need to pay attention to specific use cases, the general stereotype, which in this case is procurement of tools with embedded AI. The focus areas in the stereotype, give us a finite list of considerations to manage and comment

on in our final report

The AI You Didn't Know You Had

Procurement



Use Case 1:

Purchase and use of Third-party systems with embedded AI - Many and varied uses of AI.

eg: *Adobe Acrobat, Microsoft Office, etc*

Stereotype:

Third-party tool vendor risk management

Jurisdictions:

England & Wales

Focus:

- Address specific rogue unacceptable use cases
- Add any contractual terms needed to govern AI to existing MSAs, SOWs, etc
- Adjust procurement policy
- Introduce minimum sufficient 'safeguard' Policy to manage detected Sizewell C AI usage.
- Improve AI literacy.
- Reduce existing liability
- Ensure removal of prohibited use cases detected using EU AI Act terminology. bigspark validates against EU definitions of AI risk in all jurisdictions.
- Introduce inventory, test, audit and active Governance of detected High risk AI

Assumed Data: (Tick all that apply)

- Public Data
- Internal Data
- Confidential Data
- Restricted Data
- Private Data
- Critical Data
- Regulatory Data

User Groups: (Tick all that apply)

- Secure.users.Administrators
- Authorised users.Administrators
- Authorised users.Legal
- Authorised users.Auditors
- Authorised users.Employees
- Authorised users.Suppliers
- Authorised users.Management
- Authorised users.Temporary assignees
- Guest.Businesses
- Guest. Legal
- Guest.Individuals
- Other

Addressing Current AI Adoption - Sizewell C | 4

- By way of further illustration this is a list of Sizewell C's use cases:


mapped to stereotypes.

<https://docs.google.com/presentation/d/1RiU8vNLUxoe0BC7btwzC4ig3X9>

[lj2mT-wtsqulxAAnQ/edit#slide=id.g31e7c853bfc_0_79](https://docs.google.com/presentation/d/1RiU8vNLUxoe0BC7btwzC4ig3X9lj2mT-wtsqulxAAnQ/edit#slide=id.g31e7c853bfc_0_79)

- This is a rustic example of how Sizewell C's use cases can be mapped to quadrants. (note 12,3 should be positioned over the

origin and relevant in all quadrants ABC and D.



		Quadrant A	Secure	Quadrant B			
			3				
			2				
Open Data	3	2	1	1	2	3	Highly classified Data
			1				
			2				
			3				
		Quadrant C	Public	Quadrant D			

6. **Context specific question generation:** Based on the use case mapping conducted and the View 1 and view 2 sources then bigspark's human AI advisors will curate a more context specific set of questions related to the nuances of each use case as placed on the quadrants. This list of questions will be added to those derived from law, use case by use case. The resulting list of questions is now ready to give to the client stakeholders.

7. **Issue customised questionnaire to client stakeholders:** We tell our clients that we need a minimum of 4 weeks of interviews (at least) from the date a client signs up for an intention workshop – executed prior to the intention workshop and that the purpose of this period is to gather data. **Using a Microsoft Private Offering – note I see this as the Responsible AI companion app now updated with new additional content – so, the first time the client gets this we're asking for high level inputs...then the second time they see it it's updated /refreshed with specific questions for them based on those inputs.** The questionnaire is to be addressed to each stakeholder identified as being a part of the client's AI adoption program. The questionnaire includes a general section is created to capture generic 'pan – stakeholder' answers such as contact information, stakeholder names, holding company address, global risk reward stateemnts and the use case list drawn from the SOW agreement etc. **As questions are targeted at each stakeholder then, typically any one client will engage eight to ten people in submitting Q&A. To accelerate and simplify data capture, bigspark has created a bigspark Responsible AI companion app – which is to be launched in the Microsoft Marketplace, to enable the stakeholder to enter their answers. The data collected should be posted to the View 2 tables of the bigspark knowledge base when received.**

Each stakeholder will receive a set of questions tailored to their role. As an example, the finance officer might be asked a dominant set of questions about total cost of AI ownership in addition to questions about perceivef risk and reward, whereas a client legal counsel might be asked more questions about stakeholders. Each stakeholder can submit their response anonymously, or can elect to be named. By default they are anonymous.

Each stakeholder is to answer the questions as provided by the bigspark advisor, prior to the workshop. Answers are to be completed at least two weeks prior to the workshop so ideally the RAI companion app includes

notifications to remind the stakeholder to complete their submission. All answers provided are compiled as a set of raw responses. The shorthand for all answers curated for this client is View 2. View 2 contains sensitive information and so should be heavily protected. Eg. Encryption at rest etc.

8. **Hypothesis Review and Workshop Agenda creation.** After all stakeholder responses have been received but prior to the workshop, then bigspark generates a report based on **hypothesis (again this could be content refresh into the app)** – this report is synthesised at the use case level and the 2 + 10 will be examined as appropriate for each use case. For example, an analysis of the sale of an open free data organisation using AI to promote and sell chocolate cookies once a year may have relatively high advisory guidance if any in the fairness, safety and transparency sections. By contrast a use case in secret protection for a nuclear industry will have more intense and typically legal and regulated guidance on each the 2 + 10 approaches required. The report produced for the client is known as (View 3). View 3 is to be developed using (i) The bigspark knowledge base – view 1 & view 2 facts and findings (ii) Retrieval augmentation of the view 1 + view 2 sources and through human oversight.
9. **Workshop Day** - on the day of the workshop findings are reviewed and discussed with the client. Here's the agenda. It will require maximum one day.
 - Risk/Reward aspiration discussion
 - Use case review
 - Introduce the quadrant risk model and how we have mapped the use cases to the quadrant - AKA our hypothesis
 - Discuss and discover feedback
10. Report out – Based on the Intention workshop discussion, initial preparation, hypothesis test, we generate a report. We have termed this report output View 3. View 3 will be given to the client. Must be professional and, might be viewed while in the progress of development (TBD). Ideally much of this report generation will be automated. To start it might be manual (TBD). Dimple to look at getting AI assistance through RAG techniques.

Accelerators (Step 1)

Note – The item in Red is the MVP item I think we need to get into the Marketplace as a public offering to serve Sizewell C. We should attempt to develop this minimalistic app next week.

Audiences

There are two primary product audiences for the tools we develop. A list of audiences and tools for the AI Advisory version 1.0 tool suite are listed below.

1. Bigspark AI advisors – bigspark AI advisors are always employees of bigspark
 - a. **Data Entry: Bigspark employees enter , legal, technical and design information** to build a **knowledge base**. The scope of the knowledge base is driven by curating laws, regulations, policies, and specifications and then asserting the requirements as rules mapped to principles. As we only have 2+ 10 principles this job, while time consuming and detailed is finite. The curation of information and guidance derived from LRPS and best practice is known as ‘View 1 information’ **(This is what Maria, Abdul, Frank and now Joel have been working on)**
 - b. **Context Assessment performed in the responsible AI companion app: Bigspark advisors use a quadrant based risk assessment methodology to create stereotypes of AI use and this assessment relies on understanding the strategy context and purpose of the client’s intended and actual AI use.** Reuse of the risk assessment quadrants and stereotypes makes it possible for a lot of recommendation to be researched and prepared in advance. When first approached clients are asked to list up to 10 prioritised use cases.

When use cases are received bigspark advisors either map use cases to existing AI stereotypes, and/or create new stereotypes. Each stereotype is stored in the bigspark knowledge base.

Bigspark uses this context assessment activity to determine the price of the advisory led intention workshop. After completion of the context assessment the client can sign a SOW with bigspark Advisory. See Microsoft marketplace public offering in the client section below

c. **Intention workshop preparation. Questionnaire issuance in the RAI Companion app to all client stakeholders based on context received.** This action is done as a preparatory step for the intention workshop itself and it is undertaken for the specific client when the client has agreed scope and paid . In synopsis the bigspark advisory team create a set of more detailed Q&A to assess opportunity and risk based on context.

Our working team have termed this client specific data collection activity as 'View 2 data collection'

To achieve this data collection, a targeted set of questions are loaded into the Microsoft marketplace private offering for the client to comment on and answer – see Clients section below point b. *(Generating the questions can be done manually, by selecting a list from our knowledge base, but we've asked Dimple to look at automating more of this through RAG, assuming that the AI is on prem and custom and network isolated. Ideally I think it's automated candidate question generation with human oversight applied before submission to the client private offering. See next. I also think that the same physical app will be used in the private offering, we're simply loading additional , client specific content. Indeed the vision is that*

once we have got this app in the hands of stakeholders, it becomes the one stop shop for bigspark to add content and deliver it to the client)

- d. **Report generation.** After the Intention workshop and on the basis of information in the knowledge base a report is produced for the client. The report summarises client findings, recommendations and roadmap. This report will include a risk assessment score, and if the client has already adopted AI the roadmap will include suggested remediations. Other factors in a roadmap might include introducing governance, developing a training program and so on.

2. Clients

- a. **Public offering Microsoft marketplace.** A simple, standard questionnaire app is offered to the client at pricepoint \$0 to solicit NDA signing, filing, and to collect preparatory input from client regarding use cases (based on stereotype cards provided), risk and reward expectations (based on extensible pick lists of options) as the basis of scoping planning and pricing an intention workshop. Completion and response to the questions in this app must be undertaken before a client can have an intention workshop. **Outcome:**, Jurisdiction, industry, use case and participating stakeholder count agreed. Intention workshop scope and price agreed. Appropriate SOW/MSA signed. Use case priorities agreed.
- b. **Private offering Microsoft Marketplace Automate the Capture information by stakeholder.** After a client has agreed scope and price through the private offering we issue a curated set of questions to all stakeholders

APPENDIX –

Bigspark's Knowledge Base Accelerator specification

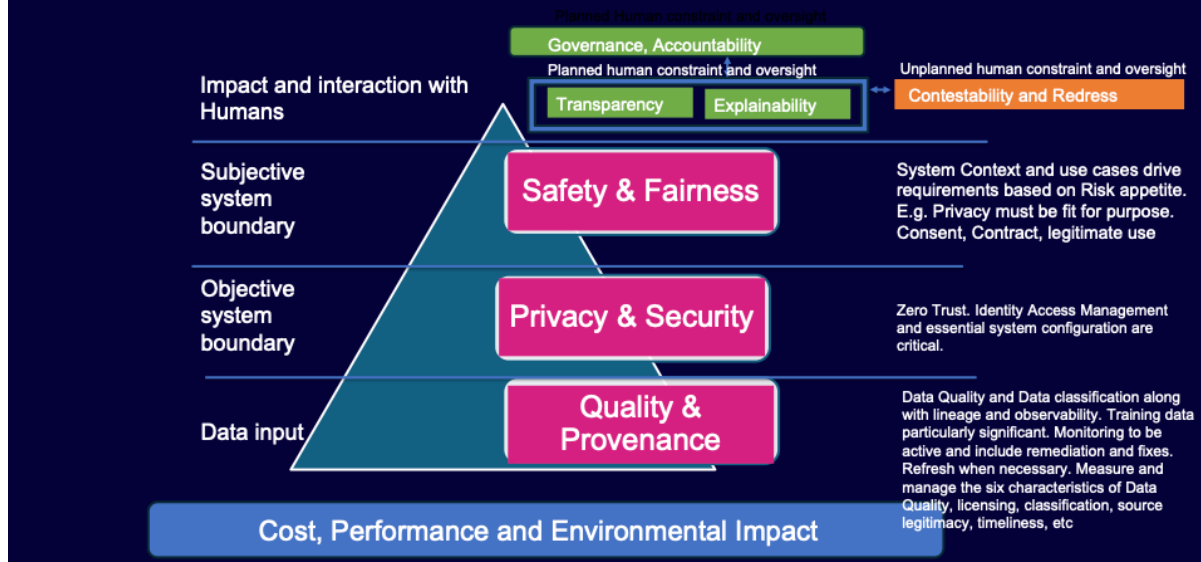
The Following content is about 3 months old, but is the original specification and set of ideas we used to start work on the Knowledge base

Goal:

The AI advisory knowledge base set is designed to assist bigspark's AI advisors in assessing and mitigating risks associated with AI systems. It provides a comprehensive framework for:

- **Bigspark AI Knowledge management:** Storing and organizing information related to laws, regulations, rules, best and customer data. This body of knowledge, such as 'bigspark questions', 'bigspark rules' etc can be curated independent of the client as they are applicable to all clients within certain parameters. As an example, all clients in the financial services industry are subject to FCA regulation. It is not client specific. As a short hand we call the sum of all information and facts in the Bigspark AI knowledge management database **View 1**. There are three truly differentiated items of knowledge in the bigspark knowledge management database.
 - 1. The 2 + 10 principles. We have Data Governance and Quality and ten attributes of responsible AI identified. Laws, regulations, policies, specifications and best practices map to rules aligned with principles. This allows us to maintain traceability from business/legal/policy through principle and rule to outcome

Our method employs a proven '2 + 10' Framework based in Law, IT and Design



-
- Bigspark Quadrant. The bigspark quadrant is an ABCD based analysis method leveraging
-
- **Client specific information:** Before advising a client specifically we ask them to supply their priority use cases and state the top three Rewards and top three Risks that they perceive are relevant to AI adoption and literacy. These questions are asked of each stakeholder so more than three risks and more than three rewards will be collated. We ask the client to summarise their top 10 standard Use cases, and, if possible, supply metrics for each use case. For a workshop with a standard price only 10 use cases are allowed. We collect the following data for each use case. 1. Data classifications. 2. User authorizations and authentications.
- **Risk assessment:** Identifying and evaluating potential risks of using AI in different applications and scenarios by clients. Bigspark assesses risk by mapping use cases to

- **Remediation planning:** Developing strategies to address identified risks, leveraging AI to suggest optimal solutions.
- **Proposal and report generation:** Creating automated documents outlining risk assessments and remediation recommendations.

User Personas

1. Data Steward (Maria, Kate, bigspark best practice leaders and founders):

Primary Goal: Ensuring data sourcing, accuracy, consistency, and compliance.

Key Activities:

Loading and managing LRPs, rules, best practices and remediations.

Creating and maintaining data augmentation tables.

Reviewing and approving rule changes.

Needs:

A user-friendly interface for data entry and management.

Clear guidelines and templates for creating and maintaining reference information.

Tools for validating data quality and consistency.

2. AI Advisor:

Primary Goal: Assessing risk and recommending remediation strategies. Managing tasks and client-related processes.

Key Activities:

Grouping rules for customer risk assessments.

Creating and managing action plans.

Conducting assessments (interviews, data analysis, etc.).

Scoring and summarizing assessment results.

Needs:

A streamlined workflow for risk assessment and remediation recommendations.

Access to relevant data and tools for conducting assessments.

Clear guidance on scoring and summarizing results.

3. Administrator:

Primary Goal: Overseeing the system and ensuring its smooth operation.

Key Activities:

Managing user accounts and permissions.

Monitoring system performance and troubleshooting issues.

Configuring system settings and parameters.

Needs:

A comprehensive system overview and management dashboard.
Tools for monitoring system performance and troubleshooting issues.
Clear documentation and training materials for system administration.

Data Entry and Reference Information (Admin User)

- Upload Laws, Policies, Regulations, and Standards (LRPS) with source references and versions.
- **Create LRP**
 - Go to LRP Addition screen
 - Create LRPs by uploading document containing LRP or
 - Create LRPs by pasting LRP contents into text window
 - Add sources and references
 - Tag as Law, Regulation, Policy or Standard
 - Tag with jurisdiction, industry, other meta data
 - Save version of LRP
- **View All LRPs and best practices**
 - See list of LRPs in system with key information displayed - Key info: Name, classification, creator, version, etc
 - Open to view single LRP.
 - See versions of LRP
- **View or Edit Individual LRP**
 - Open LRP and see versions.
 - Open version to in view mode to view
 - Open in edit mode to modify metadata, classification, naming etc
 - When modifications are made to contents of law, save as new version
- Define a classification system for LRPs (e.g., industry, jurisdiction).
- Create standard and custom LRPs with descriptions, summaries, and original source links.

- Develop a system for data augmentation tables (e.g., industry codes, vocabulary lists) for AI and administrators.

Rule Creation and Management (Admin User)

- **Create Rule**
 - Go to Rule creation screen
 - Create Rule by providing basic details - name, type, properties, classification
 - Enter rule contents
 - Set rule complexity (low, medium, high)
 - Select LRPs to link to rule
 - Save as first version
- **Update Rule**
 - See list of Rules
 - Select rule in view mode
 - Click to edit rule contents or meta information
 - Save as new version if contents of rule are changed
 - Save as new version of rule if linked LRPs are modified
- Link LRPs to a set of relevant rules.
- Classify rules as Design, Tech, Legal & Regulatory (extendable).
- Allow rule versioning with change tracking.
- Enable AI and human authored rules with separate editable sections.
- Define a rule structure including:
 - Rule number and description
 - Rule type
 - Qualification (condition to be met)
 - Assessment type (interview, test, etc.)
 - Metric (number or threshold for compliance)
 - Consequence of non-compliance

- Recommended remediation(s)
- Scoring system for validation
- Assign rules to specific LRPs and customers.

Customer Management (Admin, Advisor)

- Create customer cards with contact information and RACI structure (Responsible, Accountable, Consulted, Informed).
- **Create Customer**
 - Select create customer action
 - Enter customer information - business name, address, RACI contacts
 - Require at least 1 contact
 - Future version: Pull customer information from CRM
 - Save customer in system

Remediation Management

- Establish a system for entering and managing remediations with details and versions.
- Create groups by priority and sequences within groups by time for remediations.
- Allow grouping of remediations into categories (e.g., Security, Privacy).
- **Create Remediation**
 - Select Create Remediation from Remediations page
 -

Assignment Management

- An Assignment is the central entity that links rules to customer assessment. It lives through the lifecycle of rules being generated, customer assessments and remediation actions.

- The assignment offers a holistic view of the risk assessment process, including rules, actions, evidence, and AI insights.
- An assignment can be created from a customer page and then appropriate rule grouping done, or created directly from a set of rules and mapped to a customer.
- *All actions attached to the rules will be attached by default and can be deleted by a human user.* All actions attached to rules via conditions - conditions have actions by default.
- **Rule Grouping:**
 - An assignment is created by grouping relevant rules and assigning to a particular customer.
 - It involves selecting existing rules or authoring new ones as needed.
- **Action Generation:**
 - Once the set of rules are grouped, initial actions are system generated based on the set of rules selected.
 - Additional actions can be added by the user or suggested by the AI assistant and added when the user approves.
 - Actions will involve tasks that may be completed either manually by users or completed by AI.
 - Progress on actions are tracked as tasks are performed.
 - AI assistant will give insights into the action to be taken and contextualize the action status against the growing dataset being gathered.
- **Task Assignment:**
 - These actions are assigned to AI advisors or other team members for execution.
 - Tasks are matched to advisors based on rule type. Eg a legal task will be assigned to a legal resource.
- **Evidence Collection:**

- As actions are completed, evidence is collected to support the assessment.
- Evidence can be in various formats - recordings, document files, videos, scans, photographs, interview notes, etc.

Assessment Action Creation

- Maintain a list of assessment actions (interview, data upload, tests, etc.).
- Link actions to rules.

Proposal Generation

- Develop an AI advisor to group relevant rules for customer risk assessment plans.
- Allow selection of rules by customer, LRP, or internal policy.
- Save risk assessment plans as assignments with copies of original rules for customer-specific edits.
- Generate proposal documents with:
 - Customer logo and summary
 - Estimated duration and start date
 - Count of rules, interviews, and contacts required
 - Editable description and cost sections
- Optionally include a simple resource planning table in the proposal.
- Include Excel for inclusion in proposal (this can handle cost etc generations)
- Generate an appendix listing all relevant LRPs with references.

Risk Assessment

- Assign resources to actions within customer assignments.
- Utilize action cards to standardize information collection for each action.
- Allow AI advisors to recommend remediations on action cards.
- Implement a scoring system (quantitative or qualitative) for actions.

- Calendar integration desired*
- Standard duration for each task type is used (can be overridden by Advisor)
- Interfaces for logging assessment notes, reports, details.

Final Report Generation

- Create a pre-formatted document template for final reports.
- Include client original document sources/artefacts in final document
- Advisor can modify/delete documents from report.
- Calendar of activity performed (stakeholders engaged, summary generated for tasks and processes)
- Utilize AI for generating insights and summaries of company information, actions performed, and results.
- Include details on performed actions, remediations, and conclusions.
- Final report to include benchmark against responsible 10+2 using prior customer information. (Charts (Spider charts, graphs, etc)
- All charts will be sourced from editable data in Excel)
- Generate editable reports allowing human deletion of unwanted information.