AI Strategy for the Ultimate AI Lab

Guidehouse

Overview — Building the Ultimate AI Lab

As public and private sector leaders develop data, cloud, and Artificial Intelligence (AI) / Machine Learning (ML) strategies, the need for a unified framework has never been more apparent. These strategies naturally overlap, presenting challenges spanning strategy development, implementation, continuous delivery, and operations.

Curated, tracked, and governed data served through cloud native infrastructure fuels impactful and high value AI/ML solutions. Organizations and their respective data, cloud, and AI divisions need to coalesce around a unified vision, strategy, and execution plan. Additionally, they must bring to bear rigorous engineering practices, coupled with a scientific approach to continuously deliver solutions that shorten the **time to customer value (TTCV)** and build a sustained competitive advantage.

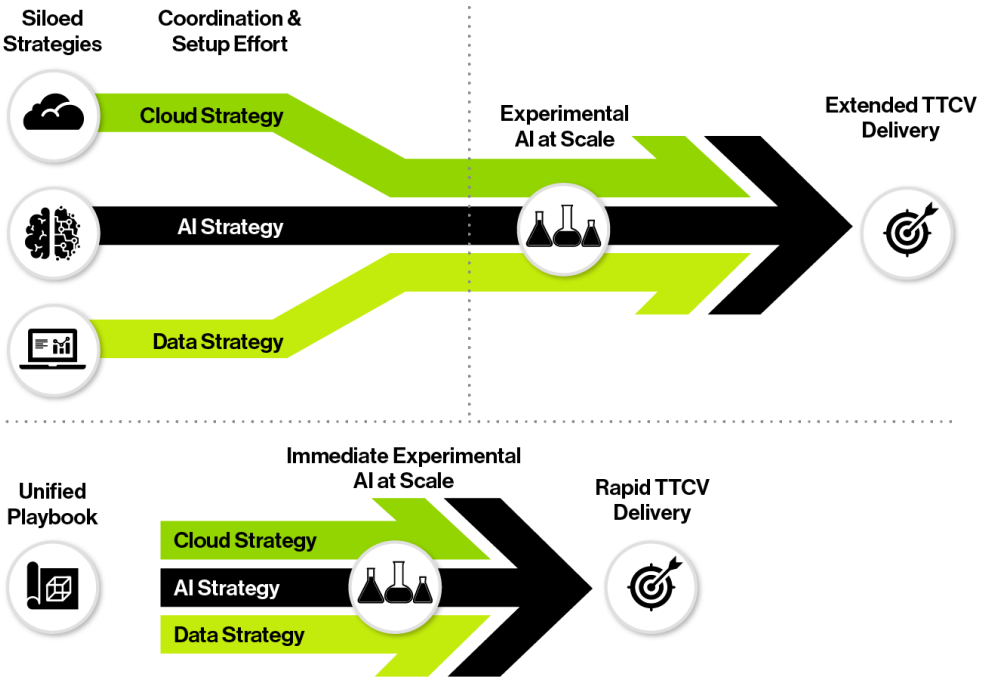

Figure 1: Unified AI, Cloud, and Data Strategies expedites TTCV

This white paper explores some of the challenges and necessary considerations when developing, governing, and releasing AI solutions. While setting an AI strategy is essential to any AI journey, it is also crucial to consider the full-stack AI ecosystem; integrated strategy enables AI teams to develop an Ultimate Lab where data is discoverable and computational resources are scalable and secure. Accordingly, this series also includes the **Cloud Strategy for the Ultimate AI Lab** and **Data Management and Governance Strategy for the Ultimate AI Lab** white papers, which explore the implementation of each respective strategy, along with a culminating business case for strategy integration: **The Ultimate AI Lab**.

AI Strategy

Adoption of AI into both day-to-day and long-term decision-making can add value and insight to an organization's existing enterprises, granting them a competitive advantage. Consequently, an increasingly common goal of executives is upscaling existing business processes with AI applications and models. Organizations can benefit from implementing AI strategy across a variety of scenarios:

- To predict impacts of different internal decisions or external factors;
- To create more efficient and accurate models and processes;
- To monitor and detect anomalies or create indicators for potential disruptions across a large amount of data; and
- To better manage and quantify risks.

Incorporating AI solutions into an organization's value proposition is an important step that requires intentional planning to fully realize its benefits. In order to yield the greatest returns to an organization's investment, its size, domain, and capacity must be accounted for within the organization's AI strategy. Integrating AI into an organization's workflow provides its own set of challenges:

- The potential for mismatch between teams and their technology; organizations that invest in AI tools and talent must ensure their teams leverage these resources.
- The potential to add high-value AI that may not reflect in business value due to lack of operational adoption; organizations seeking to operationalize AI can rely upon change management strategies to help solidify the returns on AI investments.

Change management is a critical component in maximizing the successful adoption of AI. Educating users alone falls short of that goal. At Guidehouse, we implement a five–phase framework for leading change anchored in vision, strategy, and business outcomes. It is an iterative approach to designing a curated, personal, and empowering experience for stakeholders. We employ techniques rooted in an understanding of the human psyche to create environments that nudge people toward positive outcomes. (re)Vision™ is how we create change by design.

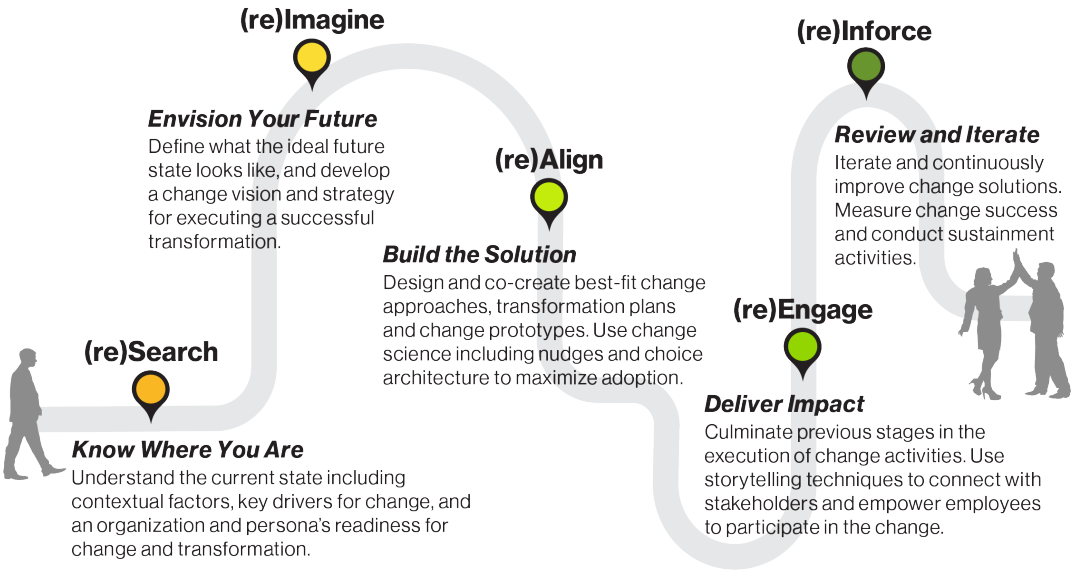

Figure 2: Guidehouse's (re)Vision™ Change Management Approach

Similarly, organizations will find that developing an AI playbook, tailored to their unique scale, requirements, and priorities will yield the strongest end-to-end value delivery.

Choosing and Adopting an AI Platform

The proliferation of AI platforms offered by cloud service providers and as commercial off-the-shelf (COTS) products presents a unique challenge to technology leaders across all organizations. In some cases, the problem is mitigated by the selection of a cloud service provider, as very few organizations would contemplate using an AI platform offered for a cloud service that they do not use. Similarly, AI platforms that do not support compliance with an important regulatory requirement can be ruled out.

Many common applications of AI have been packaged into COTS solutions or into software-as-a-service (SaaS) products. This market of COTS and SaaS AI products may serve as a viable alternative to building AI capabilities in-house. Deciding between in-house, COTS, and SaaS solutions requires technology leaders to consider what is available in the market as well as the resource constraints their organization may have.

Even after winnowing the field, technology leaders still have many options to choose from based upon their operational, staffing, security, and feature requirements. Many procurement strategies are not designed for AI products that may be repurchased many times (in the case of private data sets) or may be purchased through a consumption-based pricing scheme (in the case of SaaS offerings). To mitigate these challenges, it is critical that technology leaders set clear guidelines on the procurement of AI technology, including identifying the specific procurement vehicles that should be used.

Governing AI

CONTROLLING COSTS

AI solutions are capable of being both a cost center and a cost-control mechanism. Instead of engaging with the complexity of AI value outputs, Guidehouse recommends that organizations define value streams where AI is a constituent. Understanding how AI factors into lines of business and customer success can help AI leaders set expectations and controls for investment, unit economics, procurement, and operations.

AI SAFETY

An active subfield of AI research, "AI Safety" focuses on problems where AI must optimize for performance at a specific task without overstepping safety boundaries. This is particularly important when organizations consider using AI for mission-critical operations or when they allow AI to interact directly with employees, stakeholders, or customers. Technology leaders considering AI safety policies and safeguards must balance the need to get a return on their investment in AI with the risk profiles of the AI solutions they put into practice.[1]

EXPLAINABLE AI

In the past decade, new techniques have been developed that can enable AI systems to generate specific explanations for the decisions they make. These techniques operate based on a variety of algorithms and have been successfully integrated with AI platforms to offer on-demand explanations about AI predictions. Making AI predictions understandable can have a substantive impact on the value and governance of AI technology. In addition, explainable AI technology can lower machine and data bias, increasing accountability in the algorithm's decisions while allowing for human feedback to be incorporated into the algorithmic decision-making process. A successful AI platform gives organizations the ability to rapidly conduct AI experiments to deliver impactful insights.[2] INFERENCE LOGS

In many circumstances, AI solutions are uniquely positioned to benefit from monitoring and logging, including from a compliance standpoint. For example, triggering inference logs when poor prediction is detected or sampling predictions for review can improve model accuracy while controlling costs and time. Yet, in other instances, it may not be necessary for a firm to retain logs of the signals an AI solution consumes before making its predictions, as this approach is more costly to implement and maintain.

People and AI

PRIVACY

Stakeholders, customers, and the public have a clear expectation that AI leaders will exercise good stewardship of their private information. AI applications can create an additional burden for stewards of private information by embedding this information in opaque statistical models. AI leaders have an obligation to promote trust in society by setting clear guidelines for privacy and ensuring that they are adhered to across the organization.[3]

FAIRNESS

Organizations that make use of AI have a profound responsibility to ensure that their products do not inadvertently promote injustice. As a subfield, AI Fairness has seen robust interest in the past decade, and AI Fairness instrumentation is under active development. Technology leaders can and should leverage these opportunities to ensure fairness in their AI applications. Leaving AI to be the final word is often unethical and problematic. Including human or automated oversights in the AI development process, and appropriately using inference logs and explainability are integral parts of a fairer and more transparent AI solution.[4]

Releasing AI

FIELD EXPERIMENTS

Field experiments are often the first interaction that AI products have with the "real world." As such, AI leaders have a responsibility to create robust governance, policy, and operational support for field experiments as part of their overall AI strategy.

The same factors that make field experiments a risk element describe their value; these experiments are where AI practitioners can finally determine if their efforts to simulate the real world and real AI performance are successful. Organizations that learn how to recognize success and expand on it as part of a field test can capture value from AI.

TELEMETRY AND ACTIVE LEARNING

Impactful AI solutions have the potential to fundamentally change how their environment behaves. As a result, it is often the case that an AI model never has more value than the moment it "breaks." Models that go on accurately predicting marginal value for many years, only to fail at a critical inflection point in their environment, can be an advanced warning for decision makers.

To make AI robust and an effective bellwether, AI applications must have instrumentation for both monitoring and observability. In part to augment these measures, the emerging field of "active learning" or "augmented AI" seeks to build applications that directly support human decision-making. These technologies offer a unique value proposition to teams implementing AI in complex or regulated environments where the cost of retraining and correcting models with regular sessions of active learning may be better than risking a loss in predictive power or restarting a model from scratch.[5]

DEPLOYMENT PATTERNS

High-performing software and engineering organizations scrutinize the way that their work is deployed, and AI should be no exception. The selection of deployment patterns is of critical importance to any AI strategy, as deployment patterns impact cost, agility, security, and observability. For example, a single service inference pattern is used for light processing, while multi-model generic inference allows for hosting multiple models either as model variants or models serving different domains. Another deployment approach, microservices inference pattern, separates the steps along a deployment pipeline into distinct sub-services. This pattern can save on compute resources by isolating the complex transformations and storing their output for reuse. Technology leaders can and should align deployment patterns to operational priorities,[6] taking into account the skillset and experience of their technical resources.

Integration with Cloud and Data Strategy

EXPERIMENTS AT SCALE

Many AI experiments are resource-intensive and must leverage distributed computing clusters in order to run in a reasonable amount of time. These systems also naturally lend themselves to "big data" applications and frequently leverage cloud service providers to scale. Provisioning, administering, scaling, tuning, securing, and governing these distributed systems can often require a team of experienced technologists with diverse backgrounds (e.g., cloud engineers, database administrators, DevSecOps specialists, IT specialists, software engineers, etc.), along with effective controls and strategies to manage costs and resources spent in these efforts.

FOUNDATIONS OF QUALITY, EFFICIENCY, AND CONSISTENCY

By upholding high-quality data and high-performing architecture, organizations are preparing themselves to deliver high-quality AI solutions.

Characteristics of High-Quality Data

Curated	Data that is reliably integrated, assembled, and maintained across various sources
Data Lineage	Data that has documented, recorded, and visible origins and transformations throughout the entire lifespan of a solution
Interoperable	Data with consistent expectations for contents, context, and meaning across the systems and services where the data is created, exchanged, and consumed
Trusted	Data derived from carefully selected sources, appropriate for the intended use and audience

Amongst the qualities in the table above, organizations that develop controls to maintain data timeliness, validity, and accuracy will be prepared to deliver efficient and durable AI solutions.

Characteristics of High-Quality Solutions

Cloud	Solutions that are built and run in Amazon Web Services (AWS), Microsoft Azure, Google Cloud, Government Community Cloud (GCC) High, or other private cloud platforms
Robust	Solutions that are composed of stable, durable, end-to-end processes that produce consistent, usable output
Scalable	Solutions that can meet increased workload demands without any degradation in performance.

While the controls themselves will likely differ between various use cases, data quality assessments can provide confidence in solutions to AI practitioners and business stakeholders alike.

To promote rapid development cycles, an organization should also consider its data storage in conjunction with its AI vision. Not all data storage methods have the same strengths. For example, write-heavy databases require more compute resources and do not scale as easily as read-heavy databases. Relational databases have an inherent structure, predefined schema, vertically scalable, table based, and perform well on multi-row transactions. By contrast, a NoSQL database is non-relational, dynamic schema, horizontally scalable, and perform better on documents of JSON file formats. Additionally, NoSQL databases are document, key-value, graph, or wide-column stores.

By contextualizing data with the envisaged AI solution, organizations can better architect solutions that suit their data needs, preemptively mitigating exorbitant compute cost overruns.

Conclusion

Developing AI solutions at scale requires an AI strategy that directly incorporates data management and cloud infrastructure solutions tailored to their organization's needs. Integrating data, cloud, and AI strategies enable organizations to fully realize their investment value in each of these domains while collectively acting as a force multiplier.

A productive AI strategy enables organizations to quickly analyze and test hypotheses to produce reliable results and value-driven insights. However, if not coupled with mature data governance and a secure, agile cloud infrastructure, organizations cannot fully develop the same caliber of sophisticated, large-scale, consistent AI solutions. Segmenting, tuning, and delivering multiple AI systems necessitates the efficient allocation of cloud resources at scale and readily accessible, high-quality data. Unifying data, cloud, and AI strategies maximize return on investment while minimizing TTCV, providing organizations a competitive edge.

Guidehouse brings unique, combined expertise in applied sciences, life sciences, AI, cloud, data governance, IT strategy, and change management. This diverse but tightly coupled set of capabilities and experience uniquely position us to design and build the Ultimate AI Lab, ensuring fertile grounds for ML algorithmic experiments to discover the most promising models and guaranteeing maximum value for our customers.

[1] **Concrete Problems in AI Safety (2016)** | Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, Dan Mané

[2] **What Does Explainable AI Really Mean? A New Conceptualization of Perspectives (2017)** | Derek Doran, Sarah Schulz, Tarek R. Besold

[3] **Privacy, Algorithms, and Artificial Intelligence | The Economics of Artificial Intelligence: An Agenda (2019)** | Catherine Tucker

[4] **AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias (2018)** | Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, Yunfeng Zhang

[5] **AI-Driven Tools for Coronavirus Outbreak: Need of Active Learning and Cross-Population Train/Test Models on Multitudinal/Multimodal Data (2020)** | K. C. Santosh

[6] **Challenges in Deploying Machine Learning: a Survey of Case Studies (2021)** | Andrei Paleyes, Raoul-Gabriel Urma, Neil D. Lawrence

Guidehouse

Contact Information

Bassel Haidar
Associate Director, Advanced Analytics and Intelligent Automation
bhaidar@guidehouse.com

Charles Landau
Managing Consultant, Advanced Analytics and Intelligent Automation
clandau@guidehouse.com

Stephanie Ling
Managing Consultant, Advanced Analytics and Intelligent Automation
sling@guidehouse.com

Acknowledgments

The authors would like to acknowledge the following people for their contributions: Alex Gromadzki, Kate Pokrass, Minsoo Kim, Devan Visvalingam, and Stephen Williams

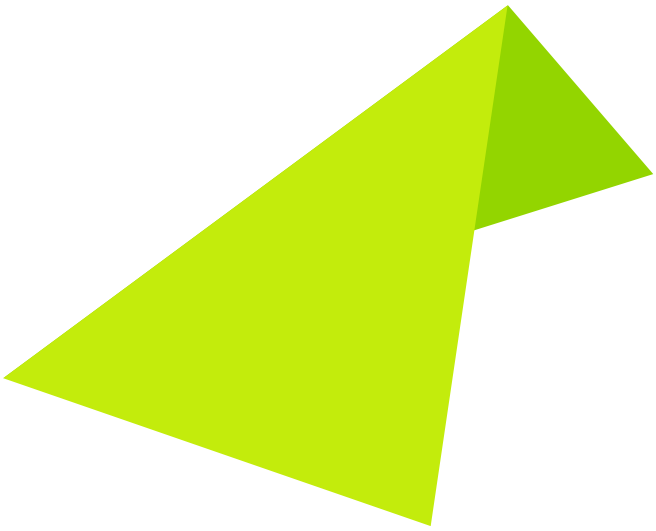

About Guidehouse

Guidehouse is a leading global provider of consulting services to the public and commercial markets, with broad capabilities in management, technology, and risk consulting. We help clients address their toughest challenges and navigate significant regulatory pressures with a focus on transformational change, business resiliency, and technology-driven innovation. Across a range of advisory, consulting, outsourcing, and digital services, we create scalable, innovative solutions that prepare our clients for future growth and success. Guidehouse is a Veritas Capital portfolio company, led by seasoned professionals with proven and diverse expertise in traditional and emerging technologies, markets, and agenda-setting issues driving national and global economies. For more information, please visit: **www.guidehouse.com**.

© 2021 Guidehouse Inc. All rights reserved. This content is for general information purposes only, and should not be used as a substitute for consultation with professional advisors. GH-125d WP AI Strategy for the Ultimate AI Lab