RIVER WEB SERVICES

A Data Center Case Study

TENASKA BUSINESS CHALLENGE

Spring Semester 2025

Proprietary and Confidential

TABLE OF CONTENTS

ATTACHMENTS

1.0 BACKGROUND

Your Company

River Web Services, Inc., (the "Company" or "River"), is a large ("hyperscale") cloud computing services provider, offering an array of products, including data storage / back-up, search, hosting for big data analytics, web applications, online game play, and access to evolving suite of proprietary generative artificial intelligence ("AI") tools. The Company delivers these services through a combination of owned and leased data centers.

The Company's AI activities include (a) the provision of customer interface for inquiries to its initial AI models, (b) the contracted development of custom AI tools and models for sophisticated corporations and governmental entities, and (c) training of new AI large language models that the Company believes will attract revenue subscriptions from individuals and smaller firms with lesser technological capability.

The Company's data centers are comprised of a large building with a specific orientation that allows (a) internal layout of racks supporting specialized state-of-the-art computer processors ("chips") and (b) necessary cooling systems to maintain the temperatures of the chips for optimal performance and lifespan.

The Current AI Market

The exact capabilities of AI technology, the use cases, and the long-term utilization of these tools by the public and private sectors is highly uncertain. Unfortunately, the only way to prove the viability of AI technology is to build the required computing infrastructure and allow the results to unfold. Thus, to a certain degree, the Company and industry are contemplating a "Field of Dreams" strategy – i.e. "build it and they (the customers and revenue) will come." There are other factors that complicate the strategic thinking:

- **Competition** – Given the potential payoff and the perception that there is a "first mover" and cost advantage to those who can scale up quickly, competition is fierce among the leading computing services providers seeking to establish a market presence.
- **Investor Sentiment** - There is intense investor pressure to pursue AI strategies even if the work to date has yet to yield a return on investment. However, it is unknown how long these shareholders will remain comfortable with the required spending given recent disappointing results of the metaverse, virtual reality, and blockchain technologies.
- **Cost** - Pursuing this "Field of Dreams" strategy will require capital investment in data center capacity that is significant in comparison to conventional cloud computing services and potentially terminal to the life of a competitor if they fail to obtain a top market presence.

On the positive side, chip technology – the backbone of AI capability – has taken a significant step forward in efficiency, such as with the NVIDIA H100 graphical processing unit ("GPU"). The advanced chips are now able to provide higher computing capability within the same sized data center or "footprint" which provides some cost and operating benefits. However, with efficiency comes increased power demand. Historically, each rack inside a data center would consume the equivalent of 2-4 kW of power (equal to 1-2 homes). However, current technology has increased to 45 kW/rack with near-commercial systems claiming 100kW/rack. As a result, energy supply, ultimately in the form of electric power, becomes the critical strategic issue for growth of data center capacity.

Leading analyst services, such as S&P Global, forecast that prospective power demand quantities ("load") for AI data centers, both in terms of individual facilities and in the aggregate, will be material to, if not beyond, the current power generation capacity and capabilities of the bulk electric transmission system (aka "grid").

Your Mandate

As leaders in the Strategic Initiatives Group, the Company's Board of Directors has charged your team with recommending a strategy to significantly expand the operating capacity of AI-capable data centers under Company control, with a primary focus on power supply, from a set of "Initial Options" (see Section 2). The recommendation must be justified against an evaluation framework of criteria established by the Board (see Section 7) and be delivered in the format of the requested Deliverables (see Section 8).

2.0 INITIAL OPTIONS

In support of this initiative, the planning team has compiled an initial set of data center build-out options to be evaluated. The scale and schedule are effectively dictated by the power supply options that are available. Following is a brief description of these options. In the Excel data file accompanying this narrative there is a table containing the key parameters for each of the options for use in the evaluation.

Contracted Power Supply Options

A. *Available Grid / Utility Supply Capacity at Industrial Rates*

   After inquiring with local utilities all over the country in locations that are suitable for data center construction, the team has determined that would be able to connect a power supply load of up to 300 MW at a location in Louisiana with a double feeder connection (two independent electric power sources) that provides 99.999% ("five nines") reliability (after accounting for the data center back-up batteries and generators).

   The power supply is immediately available from the local utilities power generation and transmission system and thus (a) the data center would simply pay the prevailing tariff (a publicly posted fixed and variable charges approved by a regulator) rate for supply and (b) start of operations by year-end 2026 would be set by the 18 month construction duration for the data center. With the fast construction schedule, River would utilize conventional cooling technology that results a relatively poor Power Usage Effectiveness ("PUE") of 1.5. The power generation would largely be from underutilized coal and natural gas capacity owned by the utility.

B. *New Larger Utility Grid Connection*

   Alternatively, another Louisiana utility is able to provide up to 600 MW via a grid connection to its system by year-end 2028. To supply the data center, the utility would consummate contracts with several renewable energy (wind and solar) and battery energy storage projects. These facilities would utilize large amounts of land. The utility would also need to build new interconnection facilities (the physical connection between to bulk power grid and the data center) and upgrade certain transmission lines. In addition, the newly dedicated transmission facilities would result in a slightly lower reliability of 99.99% ("four nines"). By building the larger data center, River would realize economies-of-scale in the capital and operating costs.

C. *New Third Party-Owned, Natural Gas-Fueled Off-Grid Power Plant*

   In order to achieve higher data center capacities, River could partner with a third-party developer to build, own, and operate a new natural gas power plant using state-of-the-art technology. By doing so, River would gain the benefits of economy-of-scale on the data center investment, economy-of-scale on the power plant, and higher fuel efficiency.

   Further, the preferred developer-partner has identified a viable location in Pennsylvania that has access to low-cost natural gas fuel and would accommodate the data center. However, selecting this option has some downsides, including: (a) without a connection to the bulk power grid, this option is forecasted to have relatively lower reliability at 99.0% ("two nines"), and (b) the soonest start of operations would be year-end 2030 due to equipment shortages.

Owned Power Generation Options

D. *Self-build a Natural Gas-Fueled Off-Grid Power Plant*

   This option is very similar to Option C above, except that River would build the power plant which would require capital investment. However, River does not possess expertise in developing, building, or operating power plants and thus, this option is expected to cost more and be less desirable than a partnership arrangement. In addition, River has only been able to identify a project site in Louisiana where current natural gas prices are higher which will increase the operating costs of the facility.

E. *Self-build a Natural Gas-Fueled Off-Grid Power Plant with Available Equipment*

Given the importance of speed to market to secure market share, River has also identified the option to invest in the construction of a smaller natural-gas fired power plant using equipment that could be supplied for start of operations by the end of 2027. However, only 300 MW of power generating capacity is available (to support a data center) with only 200 MW of IT load. In addition, the power generation equipment is costly and much less fuel efficient than other options.

F. *Acquire and Restart an Idle Coal Plant*

River has identified an available 1,500 MW coal-fueled power plant that be obtained for total of $375 MM and restarted to serve a new data center as soon as the end of 2026. Again, with the shorter schedule, River would utilize a conventional, less efficient data center configuration similar to options A and E. Further, selecting this option would result in significant water usage and air emissions. Lastly, given that this is an aged coal plant, there are high fixed operating costs (to cover maintenance) and the data center can only expect to achieve 95% power supply reliability.

3.0 <u>AI SERVICES REVENUE OUTLOOK</u>

River has engaged an outside consultant to help evaluate the overall market trends and revenue potential for the Company's products.

Revenue for AI services is expected to increase drastically in the coming years. Current estimates show that AI-related value creation could add up to $10 trillion dollars to global GDP in the next decade, both in direct revenue produced and in labor productivity gains.

Revenue forecast directly for AI tools is expected to grow at rapid pace over the 2025-2027 period, from approximately $83B in aggregate revenue in 2024 to $420B in 2027, a compound annual growth rate of approximately 72%. The growth continues after 2027 reaching $827B in 2033; however, the rate of growth between 2027 and 2030 equates to a compound average growth rate of 25%.

Due to the high projected AI revenue available, data center capacity demand is also forecast to increase dramatically. Data center capacity is expected to rise from 55 GW in 2023 to 219 GW in 2030, or a 22% compound annual growth rate.

This impressive potential growth will attract competition vying for market share. River's development team is actively working on their generative AI model. The goal is to provide customers, both individuals and institutions, a tool that can be hosted on a private cloud server to enable AI analytics on large propriety data sets while maintaining strict data privacy standards. This development effort is ongoing, and the Board has determined the related cost should be disregarded for the strategic exercise since it is an assumption common to all scenarios.

There is a trade-off between speed to market and the depth of development. This has been identified by the market consultant in both the number of customers that River can expect to acquire as well as the price that can be charged to the customers. Projects that are quicker to market will acquire more customers, and have higher initial customer growth, but will not be able to charge customers significantly more than competitors especially since this approach will offer a product that has been developed quickly and is therefore not as robust as others. Projects that are further developed will miss out on some customer acquisition and growth, but will be able to charge customers more for the tool because presumably the more time in development, the stronger their offering will be.

Information related to the market consultants revenue input is provided as Attachment 5. While the consultant has come up with a set of assumptions going into the future, these assumptions remain uncertain and an adequate risk adjusted return should be contemplated. With uncertainty around the revenue forecasts, the Board will only consider one of the options for this analysis.

4.0 VALUE OF RELIABILITY

In the near term, hyperscale data centers are anticipated to be used for training large language models. Model training applications typically require ~ 20,000 GPUs and a duration of~90 days for a data center having 200 MW of IT Load. During this period, a hyperscaler either (a) accounts for it as an internal charge among the hyperscaler's business units for it owned data centers. or (b) pays a data center lessor for the computing capacity, through a service level agreement ("SLA"). During the modeling training, this computing capacity is utilized at 100% capacity 24 hours per day.

In either case, the rates vary depending on the size of data center, technology it utilizes, location of the underlying energy source, etc. However, to quantify the value of reliability, a rate of $3.25/hr/GPU can be utilized. Assuming that a model is trained using 20,000 GPUs for 90 days at this rate, the cost basis of the model is ~$140 MM ($3.25/GPU/hr x 90 days x 24 hrs/day x 20,000 GPU's). Then, such a data center could conduct approximately four training runs each year. Larger data centers would be able to conduct additional runs for models of a similar level of computing intensity.

If a single cluster of GPUs goes down during the training run for even a few minutes or hours, the quality of the model will be reduced. For the purposes of this analysis, the Company's AI modeling team has advised you that each 0.001% of unplanned down time (during a model run) can result in a 0.005% model quality reduction (a factor of 5). This can be applied to the cost of training to get a sense of value of reliability. However, it is possible that a lower quality model will not be accepted by the customer due to inaccuracy, etc. For example, poorly trained AI models are known to "hallucinate" when providing results.

Similarly, data centers used for other purposes also place a high value on reliability given the prospective consequences of (a) interrupting critical operations of business customers and (b) consumer public outrage from loss of access to apps such as those for social media and entertainment. As such, hyperscalers typically look to build (or lease) highly-reliable data center capacity.

In evaluating the options, the Board has requested a comparison of the Net Present Value (at 15%) to be compared to the Quality Risk calculated for reliability.

Attachment 7 is an Excel spreadsheet that quantifies quality value of reliability for consideration

5.0 NATURAL GAS PRICES & VOLATILITY

For the contemplated natural gas power generation options it is important to ensure that the fuel supply is reliable while also weighing costs to operate the facility. The analyst hired to consult on natural gas fuel supply prices has provided you with an outline regarding key factors influencing pricing and volatility based on whether the facility is sited in Pennsylvania versus Louisiana. There are pros and cons related to either location based on market dynamics. Future pricing is highly uncertain and can be influenced heavily by new demand centers, changes in supply, geopolitical factors and weather.

Transco Zone 6 Non-New York (Pennsylvania)

Located in the Northeast, this pricing hub serves a region heavily influenced by weather-driven demand, particularly during winter. Proximity to the Marcellus Shale basin ensures strong supply, which can help mitigate price spikes due to local demand. However, bottlenecks in pipeline capacity during peak periods (e.g., cold winters) can lead to significant price volatility. Prices exhibit higher volatility due to extreme weather events, particularly cold snaps that drive heating demand. Pipeline constraints are a frequent issue in the Northeast, particularly during peak demand, which can lead to sharp price spikes. The market is influenced by regional demand spikes and limited interconnectivity with other hubs.

Columbia Gulf Mainline (Louisiana)

Located in the Gulf Coast, this hub is central to a vast network of pipelines and export facilities. It has more stable demand due to diverse end uses, including industrial demand, LNG exports, and proximity to Henry Hub, the benchmark for U.S. natural gas. Seasonal demand shifts are less extreme. Prices tend to be more stable due to steady industrial and export-driven demand. The region can experience some volatility from hurricanes, but infrastructure resilience and diverse supply sources typically smooth out significant price swings. The Gulf Coast benefits from a well-developed infrastructure and fewer transport bottlenecks, enabling stable supply and lower price volatility. Prices are closely aligned with Henry Hub, which serves as the benchmark for U.S. gas prices. This linkage provides a more predictable and transparent pricing environment.

Attachment 8 and the companion excel spreadsheet for this case contain forward curves of natural gas prices for these locations, including a characterization of price volatility.

6.0 FINANCING CONSIDERATIONS

As a major publicly traded computing services provider, River is sensitive to protecting its balance sheet by managing the magnitude of equity investments and liabilities. The Company has determined that target equity investment for AI data center infrastructure (including associated power generation facilities but excluding ongoing re-investment in chips over time) of $10B would not have a material negative impact on the balance sheet. Given the high risk of the investments, the Company's targeted after-tax return hurdle is >15%.

River's lenders have indicated they are willing to finance up to 80% of the cost of the buildings by a corporate term loan that is secured by the facility and backstopped by a corporate guarantee. The financing would consist of a ten-year fully amortizing term loan that carries a 7.5% rate.

The Company could theoretically increase the equity investment to $15B but would put the Company's balance sheet at risk, contingent on the AI investments panning out.

7.0 EVALUATION FRAMEWORK

The overarching objective of this strategy is to maximize the Company's prospective AI services market share, which implies (a) gaining as much market share as possible, (b) maximizing the capacity of computing resources under control, (c) achieving this as soon as practicable, and (d) achieving an appropriate risk-adjusted return. Your team must thus consider how your recommended strategy accomplishes the following:

- Controls cash flow through the amount of new data center operations (capacity and duration) to advance the AI offerings;
- Represents a cost-effective computing resource, both from economy-of-scale and marginal operating cost perspective (where energy is the primary operating cost), under the possibility that cloud computing and AI services become commoditized;
- Fits with the Company's expertise and capabilities;
- Manages the prospective impact on the balance sheet associated with capital investments; and
- May effect Company's public reputation (and customer favorability) based on the impact the data center infrastructure on land use / communities and the environmental impact associated with data center energy usage.

In addition, the team must provide an assessment of the various risks associated with the strategic options including how those risks drive analysis of the alternatives.

8.0 **DELIVERABLES**

Deliver a brief PowerPoint presentation to the Company Board that contains:

1) Recommended strategy among the options presented

2) How this recommendation fits the stated criteria

3) Summary of economic modeling results for each option that support the recommendation, utilizing the below metrics (at minimum):

  a. Total Revenue
  b. Total EBITDA
  c. Total Net Income
  d. Total Cash Flow to Equity
  e. After-Tax IRR
  f. NPV @ 15%
  g. NPV15 net of Quality Risk Measure

4) Assessment of risks associated with the recommended strategy that the Company will need to accept or seek to mitigate and how that might be accomplished

5) Basis of recommendation – how the recommend strategy compares to alternatives using the economic modeling and risk assessment

ATTACHMENT 1 – RELEVANT BACKGROUND MARKET INFORMATION

Types of Data Centers

- **AI-dedicated data centers** are an emergent class of data centers owned by either large cloud providers for their own internal AI training or inference purposes, or by third-party owners operated on behalf of hyperscaler or other corporate clients. Although very few such data centers exist currently, they are designed for the unique properties of AI workloads - high absolute power requirements, higher power density racks, and incorporate necessary infrastructure (such as liquid cooling) to accommodate these. Targeting overall power loads of 200 megawatts (MW) or higher in most cases, these facilities are typically owned by hyperscalers (Google, Amazon, Microsoft) or wholesale operators (Digital Realty, etc.)
- **Wholesale data centers** are a class of facilities that are leased to a small number of clients in relatively large blocks on a long-term basis. Wholesale data centers can be of any size but are generally larger-scale facilities 100 MW or larger with a range of either large enterprise or hyperscaler clients based on the workloads being run.
- **Hyperscaler data centers** are owned and operated by hyperscale cloud providers in which both internal (for AI, search, and e-commerce) and external (Amazon Web Services, Microsoft Azure, Google Cloud) workloads reside. Because of the scale of computing deployed for these workloads, these facilities tend to be larger in size (over 20 MW).
- **Retail colocation data centers** are facilities that are leased to a relatively larger number of clients in smaller blocks (down to several racks of equipment) on a medium- to long-term basis. Retail data centers can be of any size but are generally small- to medium-scale facilities in urban centers or nearby corporate campuses where many clients have facilities. Colocation facilities are typically situated alongside communication hubs and facilitate connectivity among clients within the same facility.
- **Corporate and telecom-owned traditional data centers** are generally owned by corporations and are often operated by third parties to accommodate internal workloads (corporate computing and storage, telecom connectivity) on a regional basis. These facilities generally tend to be small- to medium-scale in nature (under 20 MW) given their relatively localized nature.

ATTACHMENT 1 – RELEVANT BACKGROUND MARKET INFORMATION

Parameter	TIER 1	TIER 2	TIER 3	TIER 4
Uptime guarantee	99.671%	99.741%	99.982%	99.995%
Downtime per year	<28.8 hours	<22 hours	<1.6 hours	<26.3 minutes
Component redundancy	None	Partial power and cooling redundancy (partial N+1)	Full N+1	Fault tolerant (2N or 2N+1)
Concurrently maintainable	No	No	Partially	Yes
Price	$	$$	$$$	$$$$
Compartmentalization	No	No	No	Yes
Staffing	None	1 shift	1+ shift	24/7/365
Typical customer	Small companies and start-ups with simple requirements	Small & Medium Sized businesses	Growing and large businesses	Government entities and large enterprises
The main reason why companies select this tier	The most affordable data center tier	A good cost-to-performance ratio	A fine line between high performance and affordability	A fault-tolerant facility ideal for consistently high levels of traffic or processing demands

ATTACHMENT 2 – DEFINITIONS

Term	Definition
AI	Artificial Intelligence
Capacity (Power)	The maximum amount of instantaneous power (a) output that can be produced by a power plant or (b) demand for a given customer.
Hub	A contractual point (physical location) where buyers and sellers execute transactions for the purchase and sale of natural gas. Pricing among hubs varies due to differences in the supply and demand in the regional market areas.
GPU	Graphical Processor Unit; a high performance state-of-the-art computer chip or system of chips used in high intensity computing applications.
Grid	The high voltage bulk electric system comprised of transmission lines and substations over which large volumes of electrical energy are conveyed from the source of generation to large loads or local utilities for distribution to small customer loads.
Load	The amount of electricity required for a given use. When quantified on an instantaneous basis, units are usually kilowatts or Megawatts. When quantified over a period of time, units are usually kilowatt-hours or Megawatt-hours.
Maximum Power Use	The amount of instantaneous power demand for a given building when it is at 100% occupancy. Measured in kilowatts or megawatts
Megawatts	A unit of measure for power capacity. For example, a power plant may be designed for 2.5 megawatts (MW) of instantaneous output when conditions are right. 2.5 MW is equal to 2,500 kilowatts.
MW	Abbreviation for megawatts.
Megawatt-hours	A measurement of electrical energy produced or consumed. 1 MWh of electrical energy consumed in 1 hour = 1 MW of electric power.
MWh	Abbreviation for megawatt-hours.
Kilowatts	A unit of measure for power capacity. For example, a power plant may be designed for 1,000 kilowatts (kW) of instantaneous output when conditions are right.
Kilowatt-hours	A measurement of electrical energy produced or consumed.
kW	Abbreviation for kilowatts.
kWh	Abbreviation for kilowatt-hours
TWh	Terawatt-hours – A unit of measurement of electrical energy produced or consumed. 1 TWh = 1,000,000 MWh.
Emissions Rate	For this case, the amount of carbon dioxide emissions in pounds (lbs) per unit of electrical energy measured in kWh or MWh's
Power Generation	AKA electrical energy. An amount of power output over a period of time. Measured in kilowatt-hours (kWh) or megawatt-hours (MWh). For example, 1 kilowatt of power capacity that is generated for 1 hour = 1 kilowatt-hour.
Power Usage Effectiveness (PUE)	A measure of the energy efficiency of a data center, which is calculated as the total power demand / power to serve the IT Load.
IT Load	The power demand of computing equipment that is housed in a data center. Excludes cooling, lighting, etc.

ATTACHMENT 3 – FINANCIAL MODELING GUIDANCE

Operations Timing	• Once the data center and any related power infrastructure are installed, assume there is a one-year Model Training period. During this period, all expenses of running the facility are incurred, but no revenue is generated. • At the conclusion of the Model Training Period, assume customer acquisition and revenue generation begin at the beginning of the following period.
Revenue Assumptions	• Utilize the customer growth rate and pricing information contained in Attachments 5 and 6 to calculate annual revenue. • Assume that 90% subscribers generally continue to pay the annual subscription fee year after year (customer attrition rate of 10%) • Assume that information provided is “mid-year” or “annual average’ information and can be used to calculate annual revenue without any timing considerations.
Operating Costs	• All operating costs, both variable and fixed, provided are in 2025 dollars. A cost escalation of 2.5% should be applied to the base cost.
Capital Costs	• Assume both building costs and equipment costs are capitalized and depreciated. Further, assume that all costs are incurred on the last day of the period (i.e. year) and any depreciation begins in the following period. • Building Capital Costs should be straight-line amortized over 40 years. • Equipment Capital Costs should be straight-line amortized over 5 years. • Power Capital Costs should be straight-line amortized over 20 years. • Do not assume any bonus depreciation in your calculations. • All Capital Costs provided are estimated for the year placed into service, these cost should not be escalated.
Income Tax Rate	• A corporate income tax rate of 21% should be utilized in the analysis. • Assume that any tax benefits created due to negative taxable income can be utilized to offset earnings in River’s other business lines and should be included in the economics calculated (i.e., no need to calculate a net operating loss carry forward).
Debt	• As discussed in the Narrative, bank debt of up to 80% of the building capital cost can be utilized in the capital stack of the project. • The debt is assumed to be secured by the building *and* a corporate guarantee by River. • The debt offered is fully amortizing over 10 years and carries a 7.5% interest rate. • The first repayment will be in the Model Training year.
Cash Flow	• As this is a corporate project for River, assume that River covers any cash shortfall that is generated and receives any cash that is generated. • Assume cash flows through 2050 and no terminal value.

ATTACHMENT 4 – INITIAL OPTIONS DATA

Data Center Load and Cost

Power Supply Type	Contracted			Owned		
Strategy Concept Case	A	B	C	D	E	F
Name	Available Grid /Utility	Larger Grid / Utility	Gas Power Partner; Cost Optimized	Gas Power Self Build: Cost Optimized	Gas Power Self-Build: Scheduled Focused	Existing Coal Plant Restart
Data Center IT Load (MW)	200	480	1,200	1,200	200	1,000
Total Load (MW)	300	600	1,500	1,500	300	1,500
Power Usage Effectiveness	1.5	1.25	1.25	1.25	1.5	1.5
Data Center New Land Use (Acres)	100	192	300	300	100	0
Data Center Buildings Capital Cost ($MM)	$3,000	$5,760	$9,600	$9,600	$2,800	$8,000
Racks and Processors Capital Cost* ($MM)	$1,800	$4,080	$9,600	$9,600	$1,800	$8,000
Subtotal: Data Center Capital Cost ($MM)	**$4,800**	**$9,840**	**$19,200**	**$19,200**	**$4,600**	**$16,000**
Data Center Operating Costs ($/yr)	$8,000,000	$9,000,000	$12,000,000	$10,000,000	$8,000,000	$12,000,000

** Max 5 year useful life*

Power Supply, Contracting, and Pricing

Power Supply Type	Contracted			Owned		
Strategy Concept Case	A	B	C	D	E	F
Name	Available Grid /Utility	Larger Grid / Utility	Gas Power Partner; Cost Optimized	Gas Power Self Build: Cost Optimized	Gas Power Self-Build: Scheduled Focused	Existing Coal Plant Restart
Power Supply Source	Local Utility -- Exisiting Units	Local Utility - Planned Renewables and batteries	Islanded (no grid) Large efficent Units	Islanded (no grid) Large efficent Units	Islanded (no grid) small gas power units	Acquired Coal Plant
Project Location	Louisiana	Louisiana	Pennsylvania	Louisiana	Louisiana	Pennsylvania
Environmental Aspects						
Power Generation New Land Use (acres)	N/A	3350	100	100	100	0
Air Emissions (per MWh)	High	Low/No	Medium	Medium	High	Highest
Water Use	None	Low	Medium	Medium	Low	High
Schedule-Related						
Forecasted Operational Dates	End 2026	End 2028	End 2030	End 2030	End 2027	End 2026
Schedule Basis	Utility construction of interconnection	Utility Construction of Transmission and Interconnection	Equipment Lead Times	Equipment Lead Times	Permitting + Equipment Lead Times	As soon as datacenter can be built
Initial Contract Term (yrs)	1	10	15	NA	NA	NA
Power Reliability						
Net Power Reliability** (%)	99.999%	99.990%	98.000%	98.000%	99.000%	95.000%
*** After accounting for back-up batteries and generators*						
Costs						
Company Investment ($MM)	$0	$0	$0	$2,250	$750	$375
Specific Capital Cost ($/kW)	$0	$0	$0	$1,500	$2,500	$250
Contract Demand Charges ($/yr)	$18,000,000	$72,000,000	$270,000,000	$0	$0	$0
Onsite Generation Fixed Costs** ($/yr)	0	0	0	$21,600,000	$18,000,000	$54,000,000
Variable Energy Cost ($/MWh)	35	60	3	3	7	5
2026 Base Fuel Price*** ($/MMBtu)	0.00	0.00	3.37	3.87	3.87	2.50
x Heat Rate (MMBtu/MWh)	0.00	0.00	6.50	6.50	9.50	10.00
= Fuel Charge ($/MWh)	0.00	0.00	21.91	25.16	36.77	25.00
+ Variable Energy Cost ($/MWh)	35.00	60.00	3.00	3.00	7.00	5.00
= Total Variable Energy Cost ($/MWh)	35.00	60.00	24.91	28.16	43.77	30.00
x Max Power Demand (MWh/yr)	2,628,000	5,256,000	13,140,000	13,140,000	2,628,000	13,140,000
= Total Varible Energy Cost ($/yr)	91,980,000	315,360,000	327,251,700	369,956,700	115,014,420	394,200,000
+ Fixed and Demand Charges ($/yr)	18,000,000	72,000,000	270,000,000	21,600,000	18,000,000	54,000,000
= Total Operating Costs ($/yr)	109,980,000	387,360,000	597,251,700	391,556,700	133,014,420	448,200,000

**** Fuel Prices are subject to market volatility*

(Refer to separate Excel Spreadsheet)

ATTACHMENT 5 – MARKET CONSULTANT REVENUE FORECAST

Revenue Assumptions

Power Supply Type	Contracted			Owned		
Strategy Concept Case	A	B	C	D	E	F
Name	Available Grid /Utility	Larger Grid / Utility	Gas Power Partner; Cost Optimized	Gas Power Self Build: Cost Optimized	Gas Power Self-Build: Scheduled Focused	Existing Coal Plant Restart
First Year of Operations (first year of customer revenue)	2028	2030	2032	2032	2029	2028
Number of Customers in First Year	1,000,000	800,000	750,000	750,000	1,000,000	1,000,000
Price Charged / User	$140	$550	$1,500	$1,500	$160	$525

Market Growth in Customer Acquisition	
2025	72.0%
2026	72.0%
2027	72.0%
2028	26.0%
2029	26.2%
2030	23.8%
2031	20.8%
2032	17.8%
2033	14.8%
2034	11.8%
2035	8.8%
2036	5.8%
2037	5.0%
2038	5.0%
2039	4.0%
2040	4.0%
2041	3.0%
2042	3.0%
2043	3.0%
2044	3.0%
2045	3.0%
2046	3.0%
2047	3.0%
2048	3.0%
2049	3.0%
2050	3.0%

Market Solution Pricing ($/user/year)	
Chat GTP	
Plus	$ 240
Pro	$ 2,400
Microsoft Copilot	
Pro	$ 240
Google Gemini Code Assist	
Standard	$ 228
Enterprise	$ 540
GitHub Copilot	
Individual	$ 100
Organization	$ 228
Perplexity AI	
Pro	$ 240
Character.ai	
Plus	$ 120
Anthropic	
Pro	$ 216
Team	$ 300

(Refer to separate Excel Spreadsheet)

ATTACHMENT 6 – ILLUSTRATIVE EXAMPLE OF REVENUE BUILD-UP

Option A		2028	2029
Market Growth in Customer Acquisition		26.0%	26.2%
Annual New Customers		1,000,000	1,262,000
Paid Subscribers, Beginnig of Period		-	1,000,000
Less: Customer Attrition	10%	-	(100,000)
Retained Subscribers		-	900,000
New Customers Acquired		1,000,000	1,262,000
Total Paid Subscribers		1,000,000	2,162,000
Price Charged / Subscriber		$ 140	$ 140
Annual Revenue		$ 140,000,000	$ 302,680,000

(Refer to separate Excel Spreadsheet)

ATTACHMENT 7 – ILLUSTRATION OF RELIABILITY VALUE

Power Supply Type	Contracted			Owned		
Strategy Concept Case	A	B	C	D	E	F
Name	Available Grid /Utility	Larger Grid / Utility	Gas Power Partner; Cost Optimized	Gas Power Self Build: Cost Optimized	Gas Power Self-Build: Scheduled Focused	Existing Coal Plant Restart
IT Load (MW)	200	480	1200	1200	200	1000
Reliability	99.999%	99.990%	98.000%	98.000%	99.000%	95.000%
Unreliability	0.001%	0.010%	2.000%	2.000%	1.000%	5.000%
x5 = Prospective Quality Loss	0.005%	0.050%	10.000%	10.000%	5.000%	25.000%
GPUs (#)	20,000	48,000	120,000	120,000	20,000	100,000
Training Duration (d)	90	37.5	15	15	90	18
Training Run Resource (GPU*d)	1,800,000	1,800,000	1,800,000	1,800,000	1,800,000	1,800,000
Example Cost Rate ($/GPU/hr)	3.25	3.25	3.25	3.25	3.25	3.25
Training Run Cost	$140,400,000	$140,400,000	$140,400,000	$140,400,000	$140,400,000	$140,400,000
x Prospective Quality Loss =	0.005%	0.050%	10.000%	10.000%	5.000%	25.000%
Value of Quality Loss/Run	$7,020	$70,200	$14,040,000	$14,040,000	$7,020,000	$35,100,000
x Max Runs per year	4	9	24	24	4	20
= Quality Risk / Year	**$28,080**	**$631,800**	**$336,960,000**	**$336,960,000**	**$28,080,000**	**$702,000,000**

(Refer to separate Excel Spreadsheet)

ATTACHMENT 8 – FORWARD GAS CURVES

Calendar Year	Columbia Gulf (Louisiana)			Transco Non-New York (Pennsylvani)		
	Forward Price Forecast	Volatility As Annual Average		Forward Price Forecast	Volatility As Annual Average	
	($/MMBtu)	Min	Max	($/MMBtu)	Min	Max
2026	$3.87	95%	113%	$3.37	85%	159%
2027	$3.61	94%	114%	$3.19	84%	163%
2028	$3.36	94%	115%	$3.20	84%	163%
2029	$3.27	94%	115%	$3.12	84%	164%
2030	$3.13	94%	116%	$2.81	82%	171%
2031	$2.99	93%	117%	$2.64	81%	176%
2032	$3.01	93%	117%	$2.66	81%	175%
2033	$3.01	93%	117%	$2.62	81%	176%
2034	$3.12	94%	116%	$2.71	82%	174%
2035	$3.23	94%	115%	$2.82	82%	171%
2036	$3.03	93%	117%	$2.61	81%	177%
2037	$2.93	93%	117%	$2.52	80%	179%
2038	$2.99	93%	117%	$2.58	81%	178%
2039	$3.04	93%	116%	$2.64	81%	176%
2040	$3.10	94%	116%	$2.70	81%	174%
2041	$3.15	94%	116%	$2.75	82%	173%
2042	$3.21	94%	116%	$2.81	82%	171%
2043	$3.26	94%	115%	$2.87	83%	170%
2044	$3.31	94%	115%	$2.93	83%	168%
2045	$3.37	94%	115%	$2.98	83%	167%
2046	$3.42	94%	115%	$3.04	84%	166%
2047	$3.47	94%	114%	$3.09	84%	165%
2048	$3.53	94%	114%	$3.15	84%	163%
2049	$3.58	94%	114%	$3.21	84%	162%
2050	$3.63	94%	114%	$3.26	85%	161%

(Refer to separate Excel Spreadsheet)