

Procurement Department Solicitation Addendum

Solicitation Number: 88-260027-JE
Solicitation Description: RFI Data Lake Platforms
Opening Date and Time: May 12, 2026, at 2:00 PM ET
Addendum Number: 1
Issue Date of Addendum: May 1, 2026
Procurement Agent: Jennifer Ennis

1. Response to Vendor Questions received pursuant to specifications in the RFI. Questions and responses begin on the second page of this Addendum.

#	Vendor Question	College Response
1	What are the primary enterprise systems and data sources currently in scope for the initial phase of the Data Lake initiative (for example ERP, SIS, LMS, HR, finance, identity, security tools, etc.)?	EPR data consisting of student, finance, and HR data.
2	Does Central Piedmont currently use specific Azure-native services for data and analytics today such as Azure Data Factory, Azure Synapse, Microsoft Fabric, Power BI, or Microsoft Purview that respondents should assume will remain in place?	There is limited use of these tools and should be considered independent of this RFI
3	Should respondents assume the future-state platform must be fully Azure-native, or will Central Piedmont consider hybrid or multi-cloud architectures that integrate with the existing Azure Landing Zone and governance framework?	<p>Respondents should not assume the platform must be fully Azure-native. Central Piedmont operates a hybrid environment with governance anchored in an Azure Landing Zone.</p> <p>Hybrid and multi-cloud solutions will be considered, provided they integrate with Azure-based governance, security, and identity practices. The priority is compatibility, scalability, and alignment with the College's data governance framework—not a single cloud dependency.</p>
4	Are there any required higher education regulatory frameworks, state policies, or institutional standards that should be assumed beyond general FERPA, retention, audit, and security requirements?	<p>Respondents should assume standard higher education regulatory and security requirements, including FERPA, along with applicable data retention, audit, and cybersecurity obligations.</p> <p>In addition, solutions should be capable of aligning with commonly adopted security frameworks such as CIS Critical Security Controls v8.1 and NIST SP 800-53 Rev. 5, as well as supporting institutional governance standards for data protection, access management, and compliance.</p> <p>Vendors should design for flexibility to accommodate evolving state policies and institutional requirements, even where specific mandates are not explicitly defined upfront.</p>

5	What approximate data volumes, user concurrency levels, and retention periods should respondents assume for sizing and cost modeling purposes?	The exact volume and concurrency are unknown at this time, but solutions should be scalable to enterprise-wide use. Cost modeling is not a factor considered in an RFI
6	Does Central Piedmont currently maintain an enterprise data catalog, glossary, or governance tool that should integrate with the proposed solution?	<p>Central Piedmont maintains a limited data catalog and glossary capability; however, it is not widely adopted or well-known across the institution, and does not yet represent a fully mature enterprise solution. Foundational elements such as data definitions, domain structures, and inventory efforts are actively being developed as part of the College's Data Improvement and Governance initiatives.</p> <p>Respondents should assume that enterprise data cataloging, glossary, and governance capabilities will need to be expanded and formalized as part of the future-state environment. Proposed solutions should support or integrate with these capabilities and enable improved visibility, usability, and alignment with the College's evolving governance framework.</p>
7	Are there specific SIEM, identity, or security monitoring platforms currently in use that the solution should integrate with?	<p>Central Piedmont currently utilizes enterprise security platforms including Splunk Enterprise Security with Splunk Phantom for security monitoring and automation, along with Microsoft-based identity and security services.</p> <p>Respondents should assume integration with existing SIEM, SOAR, and identity platforms will be required. Proposed solutions should demonstrate the ability to integrate with these technologies, support centralized logging and monitoring, and align with the College's security and governance framework.</p>
8	Can CPCC describe the current state of its Azure Landing Zone — specifically which Azure regions are in use, whether a hub-and-spoke network topology is already established, and what governance guardrails (Azure Policy, Blueprints, Defender for Cloud) are currently active?	Central Piedmont has established an Azure Landing Zone as the foundation for our cloud architecture and governance. We utilize a hub-and-spoke model to support centralized services and distributed workloads, with enterprise-level governance and security

		<p>controls applied consistently across our environment.</p> <p>We are not providing specific details such as Azure regions or detailed configurations as part of this RFI. Respondents should assume that standard Azure governance guardrails are in place, including policy enforcement, security monitoring, and centralized management aligned to industry best practices.</p> <p>Azure is our primary and strategic cloud platform. Proposed solutions should demonstrate strong integration with our Azure environment, while supporting interoperability with hybrid components as needed and aligning with our enterprise controls for identity, security, networking, and resource management.</p>
9	Does CPCC currently operate any existing data lake, data warehouse, or lakehouse infrastructure — whether on-premises or in Azure — and if so, what is its current state and what migration or integration is expected?	<p>Central Piedmont does not currently operate any Data Lake. We do utilize SaaS solutions that operate their own data lakes. Demonstrating the ability of the proposed solution to interop with other data lakes would be beneficial.</p>
10	What are the primary on-premises systems that would need to integrate with the Data Lake Platform — for example, ERP, SIS, HRIS, or LMS systems — and are these systems already connected to Azure via ExpressRoute or VPN?	<p>The primary on-prem systems to Integrate with would be Ellucian Colleague using Ethos APIs. Secondary systems would include LMS (Canvas), HRIS (NeoGov), and CRM (Element 451).</p> <p>Connectivity to Azure is established via ExpressRoute with VPN as a fallback.</p>
11	Does CPCC have an existing data governance programme — including defined Data Owners, Data Stewards, and Data Custodians — and if so, what tools or processes are currently used to manage these roles?	<p>Central Piedmont has developed a data governance framework with defined roles, including Data Owners, Data Stewards, and Data Custodians; however, the framework is currently pending formal ratification.</p> <p>Foundational processes and structures are in place and being refined through ongoing Data Improvement and Governance efforts. Current management of these roles relies on a combination of documented governance practices, committee oversight, and existing enterprise tools, though maturity and adoption</p>

		<p>are still evolving toward a more centralized model.</p> <p>Respondents should assume that governance roles and responsibilities will be formally established and expanded. Proposed solutions should support role-based governance, stewardship workflows, and integration with enterprise tools, while enabling improved visibility, accountability, and scalability as the governance program matures.</p>
12	<p>What data classification framework does CPCC currently use — for example, Public / Internal / Confidential / Restricted — and is this classification currently applied consistently across all data domains?</p>	<p>Central Piedmont utilizes a data classification framework aligned to institutional needs, including categories such as Public, Internal, Restricted, and Highly Restricted.</p> <p>While this framework is defined and in use, its application is not yet fully consistent across all data domains. Adoption and enforcement are actively being expanded as part of the College’s Data Governance and Data Improvement initiatives.</p> <p>Respondents should assume that data classification is a core component of governance and will continue to mature. Proposed solutions should support classification-based controls, enable consistent application across domains, and integrate with enterprise security and data management practices to enforce appropriate protection, access, and monitoring.</p>
13	<p>Which regulatory and compliance frameworks are most critical to CPCC’s evaluation — for example, FERPA, HIPAA, GLBA, state-level North Carolina requirements, or NIST frameworks — and are there any active audit or compliance obligations that the solution must support from day one?</p>	<p>Central Piedmont considers higher education regulatory and security requirements as critical to evaluation, including FERPA, GLBA, HIPAA (where applicable), and PCI DSS, along with state-level requirements in North Carolina.</p> <p>From a security and governance perspective, alignment with recognized frameworks such as CIS Critical Security Controls v8.1 and NIST SP 800-53 Rev. 5 is expected.</p>

		We maintain ongoing audit and compliance obligations, including active audit activities. Respondents should assume that proposed solutions must support compliance from day one, including capabilities for data protection, access control, logging, monitoring, reporting, and audit readiness in alignment with our governance framework.
14	Does CPCC have an existing enterprise identity provider (such as Azure Active Directory / Entra ID) that the solution must integrate with, and are there any specific conditional access or MFA policies already in place that vendors should design around?	Central Piedmont uses Active Directory with synchronization to Entra ID for identity management. All users are required to use MFA via Microsoft Authenticator to access College resources
15	What is CPCC's preferred operating model for a Data Lake Platform — fully managed SaaS with minimal internal operational burden, a co-managed model, or a vendor-deployed solution managed internally by CPCC staff?	Central Piedmont does prefer SaaS or cloud-based solutions that minimize operational burden but it is not a deciding factor
16	Does CPCC currently use a SIEM platform (such as Microsoft Sentinel) or a security operations function that the Data Lake solution's monitoring and alerting must integrate with?	Please refer to the answer for question #7
17	What internal technical capabilities does CPCC currently have available to support a Data Lake implementation — for example, Azure engineers, data engineers, or data analysts — and what level of vendor-led implementation and knowledge transfer is anticipated?	Central Piedmont does not currently have resources, or knowledge sets to operate a Data Lake and has identified this for future upskilling and recruiting. Vendor-led implementations with knowledge transfer is the preferred implementation model.
18	What are the two or three most critical outcomes CPCC is hoping to achieve within the first twelve months of a Data Lake Platform being operational — for example, specific reporting capabilities, analytics use cases, or compliance improvements?	<p>Central Piedmont's primary focus in the first twelve months is establishing a strong, governed foundation that enables trusted, data-driven decision-making across the institution.</p> <p>The most critical outcomes include:</p> <ol style="list-style-type: none"> 1. Delivering governed, reliable reporting by leveraging validated data aligned to defined data domains, systems of record, and standardized definitions to ensure consistent, data-driven decisions. 2. Enabling prioritized analytics use cases that support student success, operational efficiency, and institutional

		<p>effectiveness through accessible, trusted data.</p> <p>3. Strengthening compliance and risk reduction through consistent application of data classification, access controls, monitoring, and audit readiness aligned to the College’s governance and cybersecurity frameworks.</p> <p>Respondents should focus on how their solution accelerates these outcomes while supporting long-term scalability for advanced analytics and AI.</p>
19	Are there incumbent vendors or existing platform investments (such as Microsoft Fabric, Databricks, Informatica, or Collibra) that respondents should be aware of or design around?	No, respondents should consider this a greenfield opportunity
20	Has CPCC established a budget range or indicative Total Cost of Ownership expectation for this initiative, and will cost model transparency be a significant evaluation factor when responses are assessed?	No budget has been allocated for this RFI. The purpose of this request is to gather market information and inform potential future procurement decisions. Any funding will be determined if and when a formal procurement is initiated. As noted in section 3.1, we are asking responders to provide an estimated total cost of ownership for the solution.
21	Will CPCC be evaluating responses against a formal scoring criteria or weighting framework, and if so, can any indication of the evaluation priorities be shared at this stage?	This is a Request for Information (RFI), not a request for services. The intent is to gather information to support potential future procurement decisions. Responses will not be evaluated for award, and no contract will be issued as a result of this RFI.
22	Can you please provide a description / architectural diagram of the “Centralized Azure Landing Zone”, as described on page 4 of your RFI?	See question 8 above
23	<p>Regarding the multiple data domains (described on page 4):</p> <p>a. Are these implemented as independent individual domains with their own governance / security / access management ?</p> <p>b. Are these individual data domains implemented on-premises or in the Azure cloud or in a hybrid model?</p>	Central Piedmont’s data domains are defined to establish ownership, accountability, and governance, not as fully independent or isolated technical environments. Governance, security, and access management are applied consistently through an enterprise framework rather than being managed independently by each domain.

		<p>From a technical perspective, these domains exist across our hybrid environment, including both on premises systems and our Azure cloud platform. Respondents should assume that data is governed consistently across environments regardless of where it resides, with integration and interoperability as key expectations.</p>
24	<p>Can you please provide details about existing Data and Analytics model?</p> <ol style="list-style-type: none"> Are these models commercially available models or custom-built models? Please describe the application stack for BI & Analytics / Data Ingestion / Data Cleansing (Curation & Quality control) How many distinct data sources currently exist? Are these data sources internal or external to the college, or both ? What is your data retention policy ? 	<p>Respondents should consider this a greenfield opportunity. Existing modeling & analytic tools, data sources and retention polices are identified in previous questions above</p>
25	<p>Please provide details (description / architecture / technologies in use) for existing enterprise identity, security monitoring, and data platforms.</p>	<p>The institution operates within a hybrid environment that includes enterprise identity, security monitoring, and data platform capabilities supporting both on-premise and cloud-based systems.</p> <p>Identity services provide centralized authentication and role-based access management across institutional systems. Security monitoring capabilities support the collection and analysis of system and application activity to enable threat detection, incident response, and audit support. Data platforms consist of a mix of systems of record and analytics environments that support reporting, data integration, and emerging data use cases.</p> <p>The institution is seeking solutions that can integrate with existing enterprise services, align with established security and governance practices, and support scalable, secure, and well-governed data management capabilities.</p>

26	<p>Can you describe the existing governance practices currently in place? Such as:</p> <ol style="list-style-type: none"> Roles Tools MDM capabilities – Data catalogue and Glossary management 	<p>Central Piedmont has established foundational data governance practices that are being formalized and matured as part of ongoing Data Improvement and Governance efforts.</p> <p>Roles are defined within the governance framework and include Data Owners, Data Stewards, and Data Custodians aligned to institutional data domains. These roles are supported through governance committees and oversight structures to ensure accountability and decision-making across the data lifecycle.</p> <p>Tools and capabilities are currently a combination of existing enterprise platforms, including collaboration tools, access management systems, and reporting environments. A limited data catalog and glossary capability exists but is not yet widely adopted, and will be expanded as part of the future-state governance model.</p> <p>Master data management capabilities are currently limited and not implemented as a formal, enterprise-wide program. Foundational elements such as data definitions, domain structures, and system-of-record alignment are being developed, with the expectation of maturing toward more centralized and governed master data practices.</p> <p>Respondents should assume an evolving governance environment and should demonstrate how their solution supports role-based governance, metadata and glossary management, and the maturation of enterprise data management capabilities.</p>
27	<p>If Central Piedmont CC decides to issue an RFP for the procurement of a solution, will vendors who did not respond to the RFI be eligible, or will eligibility be limited to those who did respond?</p>	<p>Vendors will not be limited to those who responded to the RFI. Any qualified vendor is eligible to submit a response to the RFP, regardless of prior participation in the RFI process.</p>
28	<p>Open formats vs proprietary stacks – How important is it to Central Piedmont that analytical data be stored in open, portable formats (e.g., Parquet, open table</p>	<p>Central Piedmont’s priority is for analytical data storage to be flexible, secure, and sustainable over time. We value open and portable formats</p>

	<p>formats) versus being tied to a proprietary storage/warehouse engine?</p>	<p>where they support interoperability, long-term accessibility, and reduced dependency on a single proprietary platform.</p> <p>At the same time, we require that the solution provide appropriate security, governance, access control, encryption, and audit capabilities regardless of the storage format used. We are seeking a balanced approach that supports modern analytics while preserving data portability, protection, and future architectural flexibility.</p>
29	<p>Multi-engine / multi-tool access – Do you expect multiple processing and BI engines (e.g., SQL, Python/Spark, Power BI, other tools) to access the same shared data in the lake, or are you comfortable with separate, engine-specific copies?</p>	<p>Central Piedmont is evaluating approaches that support shared access to governed data across multiple processing and analytics tools. We recognize the value of enabling multiple engines to access a common data layer to support consistency, reduce duplication, and align with data governance objectives.</p> <p>At the same time, we are open to vendor recommendations on architecture, including when separate or engine-specific data structures may be appropriate for performance, scalability, or operational considerations. The institution is seeking guidance on best practices to balance data consistency, performance, and governance while supporting diverse analytics use cases.</p>
30	<p>Coexistence with existing BI and data platforms – Which of your current BI and data platforms are must-integrate (e.g., Power BI, Tableau, legacy warehouse), and are there any you anticipate retiring if the new platform can replace their capabilities?</p>	<p>Central Piedmont currently utilizes a mix of reporting and data platforms to support institutional needs. While integration with existing tools is important, we are intentionally seeking opportunities to consolidate platforms where appropriate to reduce complexity, improve user experience, and build greater trust in institutional data.</p> <p>We are looking for solutions that can integrate with current capabilities while providing a path toward simplification, standardization, and a more unified analytics environment, with the ability to scale and support future growth in data, users, and advanced analytics needs.</p>

31	Performance expectations at scale – Do you have performance expectations or SLAs for key workloads (e.g., dashboard refresh times, concurrent users, batch window durations) that the new platform must meet?	<p>Central Piedmont is currently evaluating performance expectations for analytics workloads and is seeking guidance from respondents on typical benchmarks and best practices for supporting dashboards, concurrent users, and data processing at scale.</p> <p>Respondents should describe how their solution supports scalability and maintains reliable, responsive performance across reporting and analytics use cases as data volumes and user demand grow.</p>
32	Cost visibility and governance – How important is it for you to have fine-grained visibility into cost by workload, team, or project, and do you plan to use any kind of chargeback/showback model?	<p>Central Piedmont is evaluating approaches to cost visibility and governance and is seeking guidance from respondents on best practices for monitoring and managing costs across workloads, teams, and projects.</p> <p>Respondents should describe how their solution provides visibility into usage and cost drivers, and outline available capabilities for cost governance, including options such as chargeback or showback. Central Piedmont is interested in flexible approaches that can scale and adapt as needs evolve.</p>
33	Enterprise catalog expectations – Beyond simple technical metadata, what are your expectations for an enterprise data catalog (e.g., business glossary, certified datasets, data sharing across domains, data marketplace-style experience)?	<p>Central Piedmont’s expectation for an enterprise data catalog extends beyond basic technical metadata. We are looking for capabilities that support a business-aligned data environment, including a business glossary with defined and governed terms, identification of certified and trusted datasets, and clear alignment to data domains and systems of record.</p> <p>We also expect the ability to enable data sharing across domains in a controlled and governed manner, with role-based access and visibility into data usage. A data marketplace-style experience is desired to improve discoverability, usability, and adoption of institutional data assets.</p> <p>Overall, the expectation is to support trusted, governed, and accessible data that enables</p>

		consistent, data-driven decision-making across the College.
34	Granularity of access control – How fine-grained do you need access control to be (schema/table vs column vs row-level and dynamic data masking) across different data domains and personas?	<p>Central Piedmont is evaluating access control approaches that align with data governance, classification, and varying levels of data sensitivity across the institution.</p> <p>Respondents should describe how their solution supports fine-grained access control, including the ability to restrict access based on user role, data domain, and data sensitivity. This may include controls at the dataset, table, column, or row level, as well as capabilities such as data masking. Solutions should support a flexible and scalable model that can evolve with governance and security requirements over time.</p>
35	Cross-domain data sharing – How important is controlled data sharing across domains (e.g., institutional research vs academic departments vs IT/security) without physically duplicating data?	<p>Central Piedmont recognizes the importance of enabling controlled data sharing across domains to support institutional reporting, analytics, and decision-making.</p> <p>Respondents should describe how their solution supports secure and governed data sharing across different user groups and domains without requiring unnecessary data duplication. Solutions should enable consistent use of authoritative data while maintaining appropriate access controls and data protection aligned to governance requirements.</p>
36	Governance role readiness – While the RFI identifies Data Owners, Stewards, and Custodians, are these roles currently formalized and staffed across academic, administrative, and security domains, or do you anticipate the platform implementation including a phase to define and operationalize these accountabilities?	<p>Central Piedmont has identified governance roles such as Data Owners, Stewards, and Custodians as part of its Data Governance Framework; however, these roles are at varying levels of maturity across academic, administrative, and security domains.</p> <p>Respondents should describe how their solution can support the continued development and operationalization of these roles, including capabilities that reinforce accountability, workflow integration, and governance alignment as the institution matures.</p>

37	<p>Path to production for new data – In your current distributed analytics model, what is the path to production for a new dataset (from initial ingestion through quality checks, approval, and publication for wider use)?</p>	<p>Central Piedmont is evaluating approaches for managing the lifecycle of both new and existing data from ingestion through validation, governance, and publication for broader use.</p> <p>Respondents should describe how their solution supports an end-to-end path to production for datasets, including onboarding of existing data, data ingestion, quality validation, governance review, and controlled publication. Solutions should enable a structured, repeatable process that supports data trust, governance, and timely access while remaining flexible as institutional practices mature.</p>
38	<p>AI/ML maturity today – What is your current level of AI/ML usage (e.g., pilots, limited production use cases, no current use), and what type of workloads are you running or exploring (predictive models, GenAI, NLP on unstructured data, etc.)?</p>	<p>Central Piedmont is in the early stages of exploring AI and machine learning capabilities, with initial efforts focused on identifying potential use cases and evaluating opportunities across the institution.</p> <p>Respondents should describe how their solution supports a range of AI and ML workloads, including both emerging and more advanced use cases such as predictive analytics, natural language processing, and generative AI. Solutions should be flexible and scalable to support growth as institutional capabilities and use cases mature.</p>
39	<p>Priority AI use cases – Over the next 2–3 years, what are the highest-priority AI or ML use cases you envision (e.g., student success and retention, enrollment forecasting, scheduling optimization, security analytics, student-facing assistants)?</p>	<p>Central Piedmont anticipates expanding the use of AI and machine learning to support improved decision-making, operational efficiency, and enhanced services over the next several years.</p> <p>Respondents should describe how their solution enables the development, deployment, and scaling of AI and ML use cases across the institution. Solutions should support a range of capabilities and provide a foundation that can evolve as use cases are identified, prioritized, and operationalized over time.</p>
40	<p>GenAI vs traditional ML importance – How important are Generative AI capabilities (e.g., chatbots, natural</p>	<p>Central Piedmont sees value in both generative AI and traditional machine learning,</p>

	<p>language interfaces to data, content generation) versus more traditional predictive analytics and ML?</p>	<p>with the appropriate approach depending on the use case and institutional priorities.</p> <p>Respondents should describe how their solution supports both generative AI capabilities (such as natural language interaction and content generation) and traditional analytics and machine learning. Solutions should offer flexibility to support evolving priorities while maintaining strong governance, security, and data protection.</p>
41	<p>Model governance and responsible AI – What are your expectations around model governance (approvals, lineage from data to features to models, auditability, monitoring of drift and performance) and responsible/ethical AI (bias, fairness, explainability) in the higher-ed context?</p>	<p>Central Piedmont is evaluating approaches to model governance and responsible AI as part of its broader data governance and risk management efforts.</p> <p>Respondents should describe how their solution supports model governance across the lifecycle, including approvals, lineage, auditability, and monitoring of performance and drift. Respondents should also describe capabilities that support responsible AI practices such as bias, fairness, and explainability, and how these can scale as institutional practices mature.</p>
42	<p>Data protection for AI workloads – Are there specific privacy or regulatory constraints (e.g., FERPA, state law, institutional policy) that would prohibit certain categories of student or faculty data from being used in AI workloads, even in de-identified form?</p>	<p>Central Piedmont operates within a regulatory and policy environment that requires careful consideration of how institutional data is used, including within AI workloads.</p> <p>Certain categories of data, particularly those related to student education records and other sensitive institutional data, are subject to regulatory and institutional constraints, including FERPA and applicable data protection requirements. Use of this data in AI workloads must align with institutional policies governing privacy, security, and appropriate use.</p> <p>Respondents should describe how their solution supports data protection for AI workloads, including controls for data classification, access, de-identification, and appropriate use. Solutions should enable institutions to apply restrictions based on data</p>

		sensitivity and regulatory requirements while supporting responsible and governed use of AI.
43	Preferred model hosting approach – Do you prefer to bring your own models and frameworks, consume managed models/LLMs, or a combination, and how important is it that data, features, and models live on a single, unified platform?	<p>Central Piedmont has not established a preferred approach for model hosting and is seeking to understand available options and best practices.</p> <p>Respondents should describe how their solution supports multiple model strategies, including institution-developed models, managed models, or a combination of both. Responses should also address how their platform manages data, features, and models, including options for unified or integrated architectures that support flexibility, scalability, and governance as institutional needs evolve.</p>
44	Operational / data apps vs just reporting – Beyond reporting and analytics, do you anticipate building operational or student/employee-facing applications directly on top of the data lake (e.g., advisor dashboards, early-alert workflows, self-service student tools)?	<p>Central Piedmont has not established a defined approach for reporting, analytics, or operational data applications and is seeking to understand available options and best practices.</p> <p>Respondents should describe how their solution supports a range of capabilities, including reporting, analytics, and the development of operational or user-facing applications. Responses should outline how these capabilities can be delivered in a scalable, secure, and governed manner as institutional needs evolve.</p>
45	APIs and real-time access – How important is it for the platform to expose low-latency APIs or query endpoints (SQL/REST) so other applications and services can consume curated data products from the lakehouse without copying data?	<p>Low-latency APIs and query endpoints (e.g., SQL and REST) are important to the College’s architecture to enable real-time or near real-time access to curated, governed data products without creating duplicate data across systems.</p> <p>Respondents should describe how their platform supports secure and scalable API-based access, including enforcement of role-based access, data classification, and audit logging. The solution should ensure that real-time access does not bypass data governance,</p>

		system-of-record integrity, or data protection controls.
46	Reusable data products for apps – Would you like the platform to support reusable “data products” (versioned, governed views/tables) that can be safely shared across analytics, BI, and custom applications without creating new silos?	<p>Support for reusable, governed data products (e.g., versioned views/tables) is important to the College’s architecture. The College prefers enabling consistent, certified data assets that can be shared across analytics, BI, and applications without creating new data silos.</p> <p>Respondents should describe how their platform supports the creation, governance, versioning, and secure sharing of reusable data products, including enforcement of data definitions, access controls, and lifecycle management.</p>
47	Identity and policy reuse in apps – Should applications built on top of the lakehouse inherit the same governance and access policies (RBAC, row/column security, masking) defined for analytics users?	<p>The College does not currently operate a formal data lake or lakehouse and is evaluating options to establish a modern data platform. As part of this evaluation, the College prefers that applications built on top of the platform inherit consistent governance and access policies defined for analytics users. Consistent enforcement of RBAC, row- and column-level security, and data masking across all access methods is a desired capability.</p> <p>Respondents are encouraged to describe how their platform supports both data lake and lakehouse architectures, including how governance, policy enforcement, and access controls are consistently applied across analytics, APIs, and applications, and how the platform prevents policy bypass or inconsistent enforcement.</p>
48	Legacy ingestion priorities – Within your hybrid environment, which specific on-premises legacy systems (e.g., SIS, HR/Finance, LMS, security tools) represent the highest priority for the first phase of ingestion into the new data lake platform?	<p>The College has not yet established a defined priority for legacy system ingestion and is seeking guidance on recommended approaches. Core institutional systems such as the student information system (Colleague), HR and Finance systems, and the learning management system (LMS) are expected to be key considerations due to their role as authoritative data sources.</p>

		<p>Respondents are encouraged to describe how they would recommend prioritizing ingestion across these systems, including dependencies, data value, integration complexity, and alignment with governance and reporting needs.</p>
49	<p>Azure Landing Zone alignment – How strictly is your Azure Landing Zone architecture governed, and do you expect a new data/AI platform to run in its own subscription or within an existing spoke in your hub-and-spoke model?</p>	<p>The College is currently maturing its Azure environment and does not yet have a fully defined or consistently enforced Azure Landing Zone architecture. The current cloud architecture and deployment model are still being assessed as part of broader cloud and data strategy efforts.</p> <p>As part of this evaluation, the College is seeking guidance on recommended deployment models for a data and AI platform within Azure, including considerations for subscription design, network architecture, security, and governance. The College is open to different architectural approaches based on best practices and long-term sustainability.</p> <p>Respondents are encouraged to describe how their platform can support and align with a maturing Azure environment, including recommended approaches for establishing governance, policy enforcement, and secure architecture patterns.</p>
50	<p>Internal resourcing for long-term sustainability – To ensure long-term sustainability after implementation, what dedicated internal roles (e.g., data engineers, platform admins, institutional research analysts) do you expect to assign to work alongside the project team during build and transition, and then own day-to-day operations?</p>	<p>The College has not yet defined a final internal resourcing model and is seeking guidance on the roles and skill sets required to support implementation and long-term sustainability of a data platform.</p> <p>It is anticipated that a combination of IT, data, and institutional research resources will be involved during implementation and transition, with ongoing responsibilities likely including platform administration, data engineering, and analytics support.</p> <p>Respondents are encouraged to describe recommended roles, team structures, and level of effort required during both implementation and steady-state operations, including any</p>

		dependencies on vendor support or managed services.
51	Operational ownership & support expectations – For the steady-state operating model, how do you envision responsibilities being split between your internal teams and the vendor/platform (e.g., monitoring, incident response, upgrades, optimization), and are there specific support hours/SLAs you consider mandatory?	<p>The College has not yet defined a final steady-state operating model and is seeking guidance on recommended approaches for operational ownership and support. It is anticipated that responsibilities will be shared between internal teams and the vendor, with internal resources focused on governance, data management, and oversight, and the platform/vendor supporting infrastructure, maintenance, and platform operations.</p> <p>Respondents are encouraged to describe recommended responsibility models, including how monitoring, incident response, upgrades, and performance optimization are handled, as well as typical support structures, service levels, and any standard SLAs or support hour expectations.</p>
52	Is it mandatory for vendors to have higher education experience?	No
53	Is Central Piedmont looking for a custom solution, or is there a preference for COTS?	<p>Central Piedmont has not established a preference between a custom-built solution and a commercial off-the-shelf (COTS) platform and is seeking to understand the range of available options.</p> <p>Respondents should describe how their solution is delivered, including the use of configurable COTS platforms, custom components, or hybrid approaches. Responses should outline how the solution supports flexibility, scalability, and alignment with institutional governance, security, and operational requirements.</p>
54	Could Central Piedmont please confirm in which section of the response should vendors include the form located in RFI page 1?	It is acceptable for the execution page to be located at the beginning of your RFI response.
55	Are electronic signatures allowed?	Yes
56	What challenges are Central Piedmont currently facing with the existing data and analytics environment, including any performance or data quality issues?	Central Piedmont currently faces challenges related to a distributed data environment,

		<p>including data spread across multiple systems, unclear sources of truth, and inconsistent definitions. This has led to reduced trust in data and difficulty in identifying authoritative sources for reporting and decision-making.</p> <p>The current landscape also introduces complexity in accessing, integrating, and maintaining data, impacting consistency and usability across the institution. Central Piedmont is seeking to consolidate and streamline its data environment to improve data trust, clarity, and overall effectiveness of analytics.</p> <p>Respondents should describe how their solution addresses these challenges and supports a more unified, trusted, and governed data ecosystem.</p>
57	Can Central Piedmont share a high-level architecture of the current data ecosystem, including existing data warehouses, reporting systems, and whether this initiative will replace or coexist with them?	<p>The College does not currently have a fully defined or centralized data architecture and is in the process of assessing its existing data ecosystem. Current reporting and analytics are supported through a combination of source systems, departmental solutions, and limited centralized capabilities.</p> <p>At this stage, it has not been determined whether a new data platform would fully replace existing solutions or coexist with them during a transition period. The College anticipates a phased approach as part of its broader data strategy.</p> <p>Respondents are encouraged to describe recommended architecture patterns for institutions with evolving data environments, including approaches for integrating with or transitioning from existing reporting systems and establishing a scalable, governed data platform.</p>
58	What are the key source systems (e.g., ERP, SIS, LMS, CRM), and are there any third-party or vendor systems with known integration challenges?	<p>The College utilizes a range of core systems, including an ERP/SIS (Colleague), learning management system (LMS), CRM, HR and Finance systems, and various third-</p>

		<p>party/vendor applications that support institutional operations.</p> <p>At this stage, the College has not fully documented all integration points or identified specific systems with known integration challenges and is seeking vendor guidance in this area.</p> <p>Respondents are encouraged to describe typical integration approaches, common challenges when integrating with higher education systems (including Colleague and LMS platforms), and recommended strategies for managing third-party and vendor system integrations.</p>
59	What tools are currently used for data integration, ETL/ELT, data warehousing, reporting, and BI?	<p>Central Piedmont currently utilizes a mix of platforms to support data integration, storage, and analytics. This includes enterprise systems of record, integration and transformation processes, cloud-based data services, and reporting and business intelligence tools used across academic and administrative units.</p> <p>These capabilities support a range of operational and analytical needs but are distributed across multiple tools and approaches. As part of this initiative, the institution is seeking to evaluate opportunities to better integrate, streamline, and potentially consolidate these capabilities into a more unified, scalable, and governed data environment.</p> <p>Respondents should describe how their solution can integrate with or replace existing capabilities and support a transition toward a more simplified and cohesive data architecture.</p>
60	Which Azure services are currently in use, and how is the Azure Landing Zone structured in terms of subscriptions, environments, and governance?	<p>Azure services currently in use include foundational platform services such as Microsoft Entra ID, Azure networking, Azure Policy, and centralized monitoring and logging via Azure Monitor and Log Analytics, along with approved workload services such as Azure App Services, Storage, and data platform services where required.</p>

		<p>The Azure Landing Zone is structured using a subscription-oriented model that separates shared platform services and workload subscriptions, with further segmentation by environment (e.g., non-production and production). Governance is enforced centrally through management groups, standardized RBAC, Azure Policy, and shared security and networking services, providing consistent controls, compliance, and cost management.</p>
61	<p>Are there any constraints or preferences regarding SaaS, PaaS, or IaaS solutions, including preferred vendors or existing enterprise agreements?</p>	<p>The College has not established strict constraints or preferences regarding SaaS, PaaS, or IaaS solutions at this time and is open to evaluating options based on capability, security, scalability, and alignment with institutional goals.</p> <p>The College utilizes a mix of enterprise technologies and cloud services and is interested in understanding how proposed solutions align with common higher education environments and integrate with existing enterprise platforms.</p> <p>Respondents are encouraged to describe deployment models, vendor dependencies, and integration approaches across a range of cloud and hybrid environments.</p>
62	<p>What is the current data volume, expected growth rate, and are there real-time or near real-time data processing requirements?</p>	<p>The exact volume, growth rate, and data processing latencies are unknown at this time. Proposed solutions should be scalable for enterprise-wide use.</p>
63	<p>Do existing ETL/ELT pipelines need to be reused or migrated as part of this initiative?</p>	<p>The College has not yet completed a full assessment of existing ETL/ELT pipelines and is in the early stages of establishing data governance. As part of this initiative, the College is seeking to better understand existing data movement and transformation processes and determine how they should be governed moving forward.</p> <p>At this time, it has not been determined whether existing pipelines will be reused, refactored, or replaced. The College is interested in approaches that support transparency, data lineage, consistency of</p>

		<p>business logic, and alignment with defined data domains and ownership.</p> <p>Respondents are encouraged to describe recommended strategies for assessing, governing, and transitioning ETL/ELT processes, including how their platform supports lineage tracking, data quality validation, and enforcement of governance controls during and after migration.</p>
64	<p>What data classification standards and data retention policies are currently defined across domains?</p>	<p>The College has established an enterprise data classification framework that is applied across key data domains. Data retention policies are currently being developed and are not yet fully defined or consistently enforced across all systems.</p> <p>As part of this initiative, the College is seeking to align platform capabilities with existing data classification standards and to support the implementation and enforcement of data retention policies as they mature.</p> <p>Respondents are encouraged to describe how their platform supports enforcement of data classification and retention policies across domains, including capabilities for policy-based access control, data lifecycle management, monitoring, and alignment with governance practices.</p> <p>As a state entity, the College does comply with the North Carolina Department of Natural and Cultural Resources Records and Retention Policy.</p>
65	<p>What tools and processes are in place for data governance, including cataloging, lineage, and overall governance maturity?</p>	<p>The College is in the early stages of establishing enterprise data governance and is actively defining its governance framework, roles, and processes. Foundational elements such as data classification have been established; however, capabilities related to data cataloging, lineage, and enterprise-wide governance tooling are still maturing and not yet fully implemented.</p> <p>At present, governance practices are a combination of emerging centralized efforts</p>

		<p>and existing system- or department-level processes.</p> <p>Respondents are encouraged to describe how their platform supports data governance capabilities, including data cataloging, lineage tracking, business glossary management, and overall governance maturity, as well as how these capabilities can be implemented and scaled within an evolving governance environment.</p>
66	<p>What security and monitoring tools (e.g., SIEM) are in place, and what is the current backup and disaster recovery strategy?</p>	<p>The College has established core security and monitoring capabilities, including centralized logging and monitoring through a SIEM platform (Splunk), along with supporting security tools for vulnerability management, endpoint protection, and external threat monitoring.</p> <p>Backup and disaster recovery practices are in place across key systems; however, approaches may vary by system and are being evaluated as part of broader efforts to standardize and strengthen resilience, recovery objectives, and governance.</p> <p>Respondents are encouraged to describe how their platform integrates with enterprise security and monitoring tools (including SIEM solutions), and how it supports backup, recovery, and disaster recovery capabilities aligned with institutional requirements for availability, data protection, and operational continuity.</p>
67	<p>Which BI tools are standardized, what level of self-service analytics is expected, and are there existing data models, marts, or semantic layers?</p>	<p>The College currently utilizes a variety of reporting and analytics tools across departments, and a single standardized BI platform has not yet been established. As part of this initiative, the College is looking to move toward greater consistency in tools and reporting while still supporting self-service analytics.</p> <p>Self-service analytics is an expected capability, but it must operate within a governed framework to ensure data quality, consistency,</p>

		<p>and alignment with defined data domains and institutional standards.</p> <p>Existing data models, data marts, and semantic layers may exist within individual systems or departments; however, they are not centrally managed or standardized at this time.</p> <p>Respondents are encouraged to describe how their platform supports standardized BI capabilities, governed self-service analytics, and the development and management of shared data models, data marts, and semantic layers within an evolving data environment.</p>
68	Are there any planned or existing AI/ML or advanced analytics use cases?	<p>The College is in the early stages of exploring AI/ML and advanced analytics capabilities and has not yet defined a formal set of use cases. Initial areas of interest include supporting student success, improving operational efficiency, and enhancing data-driven decision-making.</p> <p>At this time, AI/ML efforts are exploratory and not yet standardized or governed at an enterprise level.</p> <p>Respondents are encouraged to describe how their platform supports the development and scaling of AI/ML and advanced analytics use cases within an evolving data and governance environment, including capabilities for data preparation, model development, and integration with institutional data sources.</p>
69	What monitoring and alerting tools are currently used, and what are the expected SLAs for data availability and pipeline performance?	<p>The College utilizes enterprise monitoring and alerting capabilities; however, monitoring specific to data pipelines, data availability, and platform performance is not yet fully defined or standardized and will be established as part of this initiative.</p> <p>Service level expectations for data availability and pipeline performance have not yet been formally defined. The College anticipates establishing SLAs aligned to data criticality and institutional reporting needs.</p>

		<p>Respondents are encouraged to describe how their platform supports monitoring and alerting for data pipelines and data availability, including recommended approaches for defining SLAs, performance thresholds, and visibility into data platform performance.</p>
70	<p>What is the expected number of users and concurrent workloads, and how is support currently handled (internal teams vs vendors)?</p>	<p>The College has not yet defined expected user volumes, concurrency levels, or workload patterns, as this initiative is part of a broader effort to better understand and mature institutional data usage.</p> <p>Support models are also evolving, and a formal approach to internal versus vendor-supported operations has not yet been established.</p> <p>Respondents are encouraged to describe how their platform scales to support varying levels of users and concurrent workloads, as well as recommended support models, including considerations for internal resourcing, vendor support, and long-term sustainability.</p>
71	<p>Are there budget considerations, cost optimization requirements, or pricing model preferences (consumption-based vs fixed)?</p>	<p>Budget, funding sources, and timing for this initiative are currently being evaluated as part of broader institutional planning. At this stage, the College has not established specific cost constraints or pricing model preferences.</p> <p>The College is interested in understanding different pricing approaches, including consumption-based and fixed models, and how platforms support cost transparency, forecasting, and optimization over time.</p> <p>Respondents are encouraged to describe pricing structures, cost drivers, and recommended approaches for managing and optimizing costs within an evolving data environment.</p>
72	<p>Is Central Piedmont planning to migrate existing platforms or build a new solution from scratch, and what is the expected implementation timeline?</p>	<p>The College is actively evaluating whether this initiative will involve migrating existing platforms or implementing a new solution. As part of this effort, the College is assessing its current data environment and defining a future-</p>

		<p>state architecture that supports scalability, governance, and institutional priorities.</p> <p>The implementation approach, including considerations for coexistence with existing solutions or phased transition, is being developed as part of this evaluation.</p> <p>An implementation timeline is also being defined and is expected to follow a phased approach aligned with data governance maturity, institutional priorities, and resource availability.</p> <p>Respondents are encouraged to describe recommended implementation approaches, including considerations for greenfield deployments versus migration strategies, as well as typical timelines and key phases for institutions at a similar stage of data maturity.</p>
73	<p>What level of internal expertise exists in Azure and data platforms, and what enablement or support may be required?</p>	<p>The College has established expertise in core IT and cybersecurity capabilities, including experience with cloud technologies such as Azure; however, expertise specific to enterprise data platforms, advanced analytics, and large-scale data architecture is continuing to develop.</p> <p>As part of this initiative, the College anticipates the need for enablement and support to build internal capabilities, including knowledge transfer, training, and guidance on best practices for implementing and operating a modern data platform.</p> <p>Respondents are encouraged to describe recommended enablement approaches, including training, documentation, and ongoing support models that can assist in developing internal expertise and ensuring long-term sustainability.</p>
74	<p>Provide examples of organizations with similar scale and governance needs.'- Please clarify what all information is required to show example, for instance, just a list or detailed descriptions?</p>	<p>Examples include large public higher-education institutions, multi-campus community college systems, and public-sector or quasi-government organizations that operate regulated environments, support diverse workloads, and require strong</p>

		<p>centralized governance while enabling distributed operational teams.</p> <p>These organizations typically share characteristics such as:</p> <ul style="list-style-type: none"> • Multiple operating units or campuses • Mixed production and non-production environments • Centralized identity, security, and compliance governance • Federated IT or shared-services operating models • Regulatory and data-protection obligations <p>Comparable governance patterns include standardized cloud landing zones, subscription or account segmentation by environment and function, centralized policy enforcement, and clearly defined architectural guardrails to support scalability and compliance.</p>
75	Can vendors attach a cover letter to their response?	Yes
76	(Scope) Does the College consider modern data integration, ingestion, and data management platforms that operate alongside or atop a Data Lake storage layer (for example, ELT/CDC platforms, iPaaS, API management, and master data management) as in-scope for this RFI? Or are responses intended to be limited to Data Lake storage and analytics platforms only?	<p>The College is taking a broad, exploratory approach to this initiative and considers modern data integration, ingestion, and data management capabilities to be within scope. This includes platforms and tools that operate alongside or on top of a data platform, such as ELT/CDC solutions, iPaaS, API management, and master data management.</p> <p>The College is not limiting responses solely to data storage and analytics platforms and is interested in understanding how these capabilities may be integrated or delivered as part of a cohesive solution.</p> <p>Respondents are encouraged to describe how their offerings support end-to-end data capabilities, including integration, data movement, governance, and analytics, and how these components can work together within an evolving data environment.</p>
77	(Existing infrastructure and TCO): Section 3.1 requests an estimated total cost of ownership. Should TCO estimates account for existing institutional investments,	The College is interested in understanding total cost of ownership across multiple scenarios.

	<p>such as current cloud platform agreements and enterprise software relationships, or should responses model standalone greenfield costs?</p>	<p>Respondents are encouraged to provide TCO estimates for both standalone (greenfield) implementations and approaches that align with existing institutional investments, such as current cloud platforms and enterprise software relationships.</p> <p>Responses should clearly outline key cost drivers, underlying assumptions, and considerations for cost optimization over time, including how different deployment models may impact overall cost and scalability.</p>
78	<p>(Source systems): To tailor our response to Section 4.2 (Data Ingestion & Integration), can the College share a representative list of source system categories the Data Lake is expected to ingest from (for example, SIS, LMS, HR/Payroll, Finance, CRM, identity and access management)?</p>	<p>The College anticipates ingesting data from a range of core institutional systems that support academic, administrative, and operational functions. Representative source system categories include the student information system (Colleague), learning management systems (LMS), customer relationship management (CRM), HR and payroll systems, finance systems, and identity and access management platforms.</p> <p>The College also expects to incorporate data from additional third-party and vendor systems that support institutional operations and analytics use cases.</p> <p>Respondents are encouraged to describe how their platform supports ingestion from a diverse set of source systems, including structured and unstructured data, and approaches for integrating enterprise and third-party applications within a governed data environment.</p>
79	<p>What is the estimated start date and period of performance for cost estimation purposes?</p>	<p>The estimated start date and period of performance are currently being developed as part of broader planning and funding considerations. At this stage, the College anticipates that implementation will follow a phased approach aligned with institutional priorities and data governance maturity.</p> <p>For cost estimation purposes, respondents are encouraged to provide assumptions based on typical implementation timelines for institutions of similar size and complexity, including</p>

		phased deployment scenarios where applicable.
80	Is there an established budget range or ceiling for the anticipated RFP?	See question #20
81	Could you confirm that we are required to complete and return a signed Execution page, together with responses to the items identified in Section 4.0 of the RFI for the submission?	Yes
82	What source systems are currently in scope for the data lake (e.g., Ellucian Banner/Colleague, Workday, PeopleSoft, homegrown apps)?	<p>The College anticipates that core institutional systems will be in scope for the data platform, including the student information system (Colleague), as well as HR, finance, learning management, and CRM systems. The environment also includes a number of cloud-based (SaaS) applications that support institutional operations and analytics, which are expected to be part of the broader data ecosystem.</p> <p>Additional third-party and vendor-supported applications, as well as internally developed solutions, may also be included as scope is refined.</p> <p>The specific scope of source systems will be further defined as part of this initiative and aligned with institutional priorities and governance considerations.</p> <p>Respondents are encouraged to describe how their platform supports integration with a wide range of enterprise, SaaS, third-party, and custom applications, including approaches for onboarding and scaling data ingestion over time.</p>
83	Which Azure data services does CPCC currently use (Azure Data Factory, Synapse, Azure SQL, Power BI)?	<p>The College currently utilizes Azure for select services, with limited use of Azure-native data services (e.g., Azure Data Factory, Synapse, Azure SQL, Power BI). As part of this initiative, the College is evaluating how these services may be leveraged, expanded, or aligned to support a more structured and governed data platform.</p> <p>Respondents are encouraged to describe how their platform integrates with or complements Azure data services, as well as recommended</p>

		<p>approaches for leveraging these capabilities within an evolving data environment.</p>
84	<p>Is CPCC open to a phased implementation (e.g., 2–3 priority data domains first, then expand), or does the expected scope cover all domains at once?</p>	<p>The College is open to a phased implementation approach and anticipates that an incremental rollout aligned to priority data domains may be the most effective path forward. The specific sequencing and scope have not yet been finalized and are being evaluated as part of this initiative.</p> <p>Respondents are encouraged to describe recommended phased implementation strategies, including approaches for prioritizing data domains, managing dependencies, and scaling over time to support broader institutional adoption.</p>
85	<p>Are there existing data pipelines or ETL processes that would need to be migrated or replaced, or is this a greenfield build?</p>	<p>The College is assessing its current data environment, including existing data pipelines and ETL/ELT processes, as part of this initiative. A mix of existing integrations and data movement processes is anticipated; however, the extent to which these will be migrated, refactored, or replaced is being evaluated.</p> <p>The College is open to both modernization of existing capabilities and greenfield approaches where appropriate, with an emphasis on establishing a scalable, governed data platform.</p> <p>Respondents are encouraged to describe recommended approaches for assessing and transitioning existing pipelines, as well as considerations for greenfield implementations, including strategies that support data quality, lineage, and governance.</p>
86	<p>Has CPCC adopted Microsoft Purview or any other data governance tooling, to date? If yes, to what extent (catalog only, classification, lineage)?</p>	<p>The College has adopted Microsoft Purview and is actively evaluating its capabilities to support enterprise data governance. Initial use has focused on data classification, and the College is exploring additional capabilities such as data cataloging, lineage, and broader governance features.</p>

		<p>As part of this initiative, the College is assessing how Purview or other governance tools can support a scalable, institution-wide governance model.</p> <p>Respondents are encouraged to describe how their platform integrates with or complements Microsoft Purview and similar governance tools, including support for classification, cataloging, lineage, and policy enforcement.</p>
87	Currently, are there specific problems or pain points the users are having with data governance and/or data quality?	<p>The College is in the process of establishing enterprise data governance and improving data quality practices. As governance efforts are being formalized, opportunities have been identified to enhance data consistency, standardize definitions, improve visibility into data lineage, and reduce reliance on department-specific reporting and data movement processes.</p> <p>These efforts are focused on building a more coordinated and governed approach to data management across the institution.</p> <p>Respondents are encouraged to describe how their platform supports improving data quality, consistency, and governance, including capabilities for standardization, validation, lineage, and reducing data duplication.</p>
88	How large is the data/analytics team that would operate and maintain the platform after implementation? This determines how much automation and managed services to bake in vs. expecting hands-on-keyboard administration.	<p>The College's data and analytics responsibilities are currently distributed across multiple functional areas. As part of this initiative, the College is evaluating the appropriate operating model and level of internal resourcing to support a modern data platform.</p> <p>There is interest in solutions that provide a high degree of automation and ease of administration, while still enabling the development of internal capabilities over time.</p> <p>Respondents are encouraged to describe how their platform supports efficient operations, including automation, managed service options, and approaches that reduce</p>

		operational overhead while supporting long-term sustainability.
89	Are there existing self-service BI consumers today (Power BI, Tableau, etc.), and approximately how many end users would access the platform?	<p>The College currently has a mix of reporting and analytics users across departments, with varying levels of access to data and reporting tools. Usage patterns and the extent of self-service analytics are being evaluated as part of this initiative.</p> <p>The expected number of end users and level of adoption are also being assessed, with the goal of supporting broader access to consistent and governed data capabilities.</p> <p>Respondents are encouraged to describe how their platform supports a range of user types and adoption levels, including approaches for scaling access while maintaining governance, performance, and data consistency.</p>
90	Are there power users of the analytics tools and data that will need specific services to meet their needs? Are there specific groups that use the data the most and do they have specific requirements that may be unique to them?	<p>The College has a range of users with varying levels of engagement with data and analytics, including individuals and groups with more advanced or frequent data needs. These usage patterns and specific requirements are being further assessed as part of this initiative.</p> <p>As part of this effort, the College is interested in understanding how platforms can support different user types, including those with more advanced analytical needs, while maintaining consistency, governance, and ease of use across the broader user base.</p> <p>Respondents are encouraged to describe how their platform supports a spectrum of users, including more advanced users, and how it accommodates varying levels of complexity, access, and functionality within a governed environment.</p>
91	What are the most significant or active use cases for the data and the analytics?	<p>The College has defined a set of initial use cases to support evaluation and validation of a data and analytics platform. These use cases are being used to guide requirements, validate</p>

		<p>platform capabilities, and inform the design of a scalable and governed data environment.</p> <p>As part of this initiative, the College anticipates that additional use cases will be identified and refined over time.</p> <p>Respondents are encouraged to describe how their platform supports a range of analytics use cases, including the ability to scale and adapt as requirements evolve.</p>
92	Does the data need to be accessed outside the community college by partner organizations? What are some examples of these users or groups?	<p>The College recognizes that there may be scenarios where data is shared with external partners; however, the scope, frequency, and access patterns for such use cases are being evaluated as part of this initiative.</p> <p>Any external access to data will be governed by institutional policies, data classification standards, and applicable regulatory requirements.</p> <p>Respondents are encouraged to describe how their platform supports secure data sharing with external organizations, including capabilities for access control, data protection, and auditing within a governed data environment.</p>
93	Are there groups or users whom reside outside the North Carolina that will need to access the Data Lake that will need to be considered from a performance perspective?	<p>Access by users or partner organizations outside of North Carolina is a potential consideration, though specific scenarios and usage patterns are still being evaluated.</p> <p>Performance and latency will be important factors in the platform design to ensure reliable access for geographically distributed users.</p> <p>Respondents are encouraged to describe how their platform supports distributed access, including performance optimization, scalability, and secure data delivery across regions.</p>
94	Are there different sophistication levels of the data users (Low, Medium, High)? If so, what % fall within each category?	<p>The College has a diverse user base with varying levels of data and analytics proficiency; however, specific distributions across user</p>

		<p>sophistication levels have not been formally quantified.</p> <p>Understanding user needs and maturity levels is part of this initiative and will inform how data access, tools, and capabilities are structured to support a broad range of users.</p> <p>Respondents are encouraged to describe how their platform supports users with varying levels of sophistication, including approaches for usability, training, and progression from basic to advanced analytics within a governed environment.</p>
95	What is the approximate data volume today, and what is the projected growth over 3–5 years?	<p>The College is currently assessing its data landscape to better understand overall data volume and growth trends. At this stage, a consolidated estimate of total data volume and projected growth has not been established.</p> <p>It is anticipated that data volume will grow over time as additional systems, domains, and use cases are incorporated into a centralized data platform.</p> <p>Respondents are encouraged to describe how their platform supports scalable storage and processing, including approaches for handling growth in data volume, performance, and cost over time.</p>
96	Does CPCC expect the vendor to provide ongoing managed services, or is this a build-and-transfer engagement?	<p>Central Piedmont has not established a preferred support model at this time and is seeking to understand the range of available options.</p> <p>Respondents should describe available engagement models, including vendor-led implementation, build-and-transfer, and ongoing managed services. Responses should outline how these options can support a flexible, scalable approach as institutional needs and internal capabilities evolve.</p>
97	What was the annual spend for the previous year on this project?	<p>This Request for Information is for planning purposes only. No annual spend is available, as there has not been a contract in place for these services.</p>

98	If this is a new contract, what is the estimated annual budget?	See question #20
99	Who are the previous incumbents on this project?	This Request for Information is for planning purposes only. No incumbent, as there has not been a contract in place for these services.
100	What challenges or pain points have been experienced with the current contractor?	See question #99
101	Are you open to a hybrid delivery model with a mix of offshore and onshore resources?	This information will be provided if the College determines that issuing a Request for Proposal (RFP) is warranted
102	Work will be onsite or remote?	See question #100