OSPI DATA MODERNIZATION PROJECT PART 1: FEASIBILITY STUDY

2026

Deb Came, PhD
Assistant Superintendent of Data Analysis and Support

Ted Loran
Chief Information Officer

Prepared by: Kris Hicks-Green, Project Manager

Washington Office of Superintendent of
PUBLIC INSTRUCTION

Table of Contents

SECTION 1 – EXECUTIVE SUMMARY

Washington's education system depends on accurate, timely, and trusted data to guide decisions that affect more than one million students statewide. While the Office of Superintendent of Public Instruction (OSPI) produces consistent, high-quality results; it currently operates on legacy systems that are fragmented, costly to maintain, and unable to meet increasing expectations for secure, accessible, transparent and efficient data. This feasibility study confirms that phased data modernization, prioritizing strong governance, automation, and secure technologies, is both necessary and achievable. The proposed investment, estimated at $92 million to $131 million over seven years, would create a unified and sustainable data platform that improves reporting, strengthens equity, and accountability insights, while reducing long-term operational risks.

1.1.0 Proposed Solution

OSPI's modernization initiative aims to transform the state's education data environment into a unified, secure, and accessible ecosystem. The objective is to replace aging, siloed systems with a sustainable platform that improves data quality, accelerates information delivery, and supports data-driven decisions across Washington's education programs. Modernization will position the agency to meet statewide service delivery goals and continuous improvement in student outcomes.

The approach to modernization builds on three guiding principles:

- *Governance by Design* embeds data quality, security, and compliance into system architecture and workflows.
- *Zero-Touch Operations* automates data collection, integration, and validation.
- *Self-Service Enablement* empowers educators, program staff, and policymakers with secure real-time access to trusted data.

Implementation would follow a phased proof-of-concept strategy in key data domains and governance foundations. This approach enables OSPI to validate solutions, build staff capacity, and manage change while demonstrating value and reducing risk. Each phase will deliver incremental benefits throughout the modernization lifecycle before scaling statewide.

1.2.0 Benefits

Modernization will deliver measurable improvements in how Washington collects, manages, and uses education data. These benefits enhance statewide transparency and ensure greater value from public investment:

- A single, authoritative source for education data statewide.
- Faster and more accurate funding, accountability, and performance reporting.
- Enhanced security and privacy through automated controls.

- Provide better investment value through streamlined systems and improved service delivery.
- Expanded workforce capacity through automation and new skill development.
- Greater staff retention opportunities by modernizing workforce skills, expanding training opportunities, and adopting modern technologies.
- Improve collaboration and information sharing across districts and programs.

1.3.0 Costs

The total cost estimate incorporates technology, staffing, training, operations, and vendor partnerships with a scalable, secure foundation capable of supporting Washington's evolving education data needs.

7-Year Costs Summary (Implementation + Maintenance & Operations)			
Cost Category	Microsoft	Amazon Web Services	Google Cloud Provider
Implementation	$87.4M	$48.5M	$49.9M
Maintenance & Operations	$43.6M	$45.2M	$42.4M
7 Year Totals	**$131M**	**$93.7M**	**$92.2M**

1.4.0 Risks

Successful modernization depends on proactive management of three primary risks:

- Platform Sustainability: The chosen cloud environment could exceed support capacity. Managed by aligning platform choice to internal capabilities and defining clear ownership and hand-off plans.
- Workforce Readiness: OSPI has gaps in cloud and data engineering. Address through early hiring, vendor support, and targeted training programs.
- Governance Consistency: Policies and standards may be applied unevenly across teams. Mitigate through a centralized Data Governance Council, automation, and standardized tooling.

1.5.0 Conclusion

Modernizing OSPI's data systems is a strategic opportunity to strengthen transparency, efficiency, and equity in Washington's education system. A modern, integrated data platform will enable timely, informed decisions that improve educational outcomes and ensure the state's data infrastructure is reliable, secure, and built for the future.

SECTION 2 – BACKGROUND AND NEEDS ASSESSMENT

OSPI's current information technology and data operations environments reflect decades of incremental growth that has supported essential educational services but now shows the strain of legacy systems, fragmented governance, and limited automation. It is necessary to establish a clear understanding of the agency's existing architecture, workforce capacity, and operational processes. This section provides a foundation through a comprehensive assessment of OSPI's current state environment, identifies the conditions, challenges, and dependencies that shape future modernization efforts.

The following analysis evaluates OSPI's infrastructure, applications, tools, governance practices, and organizational readiness to determine how well current systems align with the agency's evolving mission and statutory requirements. While the technology landscape remains functional and reliable in many areas, it relies heavily on manual processes, aging platforms, and informal coordination structures that limit efficiency, scalability, and real-time insight generation.

Key themes from this assessment include:

- **Legacy complexity and integration challenges** with multiple standalone systems and data stores lack interoperability, creating inefficiencies and limiting data driven decision-making.
- **Manual and inconsistent processes** where operational workflows depend on manual data entry, reconciliation, and reporting which introduce delays and quality risks.
- **Limited automation and governance maturity** where policies and standards exist but are applied inconsistently with minimal automation to enforce data quality, access, and compliance.
- **Evolving workforce readiness** with staff maintained deep institutional knowledge, but the need for growth and expanded skills and capacity to support modern, cloud-based operations.
- **Organizational and change readiness** with clear communication, defined roles, and leadership alignment are critical to ensure successful adoption of future state capabilities.

This baseline analysis defines the technical and organizational realities that modernization must address. The subsections that follow present detailed findings across infrastructure, applications, processes, governance, and workforce dimensions, forming the evidence base for the recommendations and future-state design outlined in later sections.

Understanding OSPI's current business environment is essential to identifying where current strengths can be built upon and where modernization will deliver the greatest impact. Technology, people, and process play a critical role in how effectively the agency can collect, manage, and use data. While OSPI's systems and staff have supported statewide programs

reliably for years, the current environment shows growing strain from legacy platforms, limited automation, and manual coordination across teams. Examining these three areas provides a clearer view of the everyday challenges and dependencies that modernization must address to build a more integrated, agile, and data-driven organization.

2.1.1. Technology

OSPI's technology foundation supports a broad range of mission-critical systems that enable statewide education programs, fiscal operations, and compliance reporting. While these systems have proven dependable, most remain rooted in traditional on-premises architectures and rely on manual processes for integration and maintenance. The result is a technology environment that is relatively stable but increasingly difficult to scale, automate, or adapt to new policy and data needs. Understanding this landscape is essential to defining the modernization path forward.

Infrastructure

OSPI's infrastructure environment supports a mix of core business applications, data services, and end-user tools at present. The technology landscape is hosted primarily in traditional environments, which include on-premises and co-location facilities in the Washington State Data Center, with limited adoption of public cloud services to date.

Key infrastructure domains include:

- Hosting is predominantly on-premises and in co-location facilities in Washington State Data Center.
- Network topology includes connectivity between agency offices, districts, and external partners.
- Identity & Access Management (IAM) across OSPI systems varies. Foundationally, the custom-built Education Data System (EDS) is the main authentication portal for determining identity and access, with some tools leveraging Entra ID.
- Security controls feature firewall protections and endpoint security. Key controls, including role-based access controls (RBAC), conditional access, encryption policies, and key management vary in maturity over time and may impact OSPI's future state adoption of zero-trust principles and maintain visibility as systems migrate to cloud environments.
- Business continuity and disaster recovery are reliant on legacy setups, and current systems may need to undergo regular backup validation or recovery testing.
- Infrastructure provisioning and management show no clear Infrastructure-as-Code (IaC) strategy in tools like Bicep or Terraform.
- System monitoring appears to be relying on built-in monitoring from each tool without a clear strategy for centralization.

Gap Notes: Comprehensive diagrams and configuration documentation for hosting, networking, IAM, and security controls varies and have clear gaps from system-to-system.

Applications

OSPI operates a comprehensive portfolio of 42+ web applications supporting statewide educational operations, reporting, and compliance requirements. Infrastructure discovery has revealed a substantial technology footprint built primarily on Microsoft .NET/IIS architecture with backend SQL Server databases.

Discovered Application Landscape

Applications seem to be a varying mix of on-prem .NET Microsoft applications, to cloud Power Apps, to other SaaS solutions that are all accessible through EDS.

- Core educational systems consist of applications that run through EDS with development, test, and production environments supporting student information management and academic data workflows. Most EDS applications operate on a legacy framework that may need upgrading to the cloud.
- Financial management includes a dedicated Financial Management System with production and development instances managing budget tracking, fund distribution, and financial reporting. Interviews mention Salesforce as the new Grant Management application.
- Reporting and analytics encompass Report Card applications and Report Card Tableau systems providing reporting and data visualization capabilities for educational performance tracking to the public.
- Specialized applications cover domain-specific systems, including Traffic Safety education management, Apportionment Services for fund distribution calculations, and various administrative consoles.

Infrastructure Foundation

- The infrastructure foundation features standardization on .NET/IIS applications hosted on Windows Server infrastructure.
- The database layer utilizes SQL Server database instances with several user databases per instance, indicating significant data complexity.
- The server infrastructure is comprised of over a hundred total servers, both Windows and Linux, supporting the application portfolio.
- Environment management maintains a separation framework of production, development, and test environments across key systems.

Key Considerations

- Integration patterns show prevalent direct database connections for production work with limited-service layer framework ensuring governance.
- Documentation gaps exist especially with application dependencies, data flows, and integration mappings; data lineage documentation is manual and thus not fully documented for all data flows.

- Deployment maturity indicates traditional deployment approaches using Team Foundation Server (TFS), with projects already underway for Continuous Integration/Continuous Delivery (CI/CD) and Development, Security, and Operations (DevSecOps) modernization in Azure Development and Operations (DevOps).

Tools Inventory

OSPI's current tooling environment reflects a mix of traditional on-premises utilities, Microsoft platforms, and limited public cloud adoption. While some enterprise tools like Power BI and Power Apps have emerged in pockets, the broader OSPI tools ecosystem lacks centralized oversight, tool standardization, or a defined tooling strategy. The following table outlines known tools and platforms in use as discovered during the infrastructure and application review:

Discovered Tooling Landscape

Category	Tool	Description
Cloud Platforms	Microsoft Power Platform	Primary cloud platform; used for app hosting and analytics (e.g., Power Apps, Power BI)
	Microsoft Fabric	IT and Data teams are using tools in Azure like Fabric to explore better analytics capabilities; no wide adoption noted, Child Nutrition team has an instance of Fabric in production-level, HR is working on standing up an instance for production.
Applications	.NET / C# / ASP.NET	On-prem application suite under Microsoft that runs about 60% of all OSPI apps interfaced through EDS.
	JavaScript	Used for frontend development in Information and Condition of Schools (ICOS) to enable user interactions like filtering and downloading reports in PDF, Excel, or Word formats.
	ESRI GIS Server and software	ESRI GIS Server used to store Geographic spatial data pertaining to calculating and optimizing School Bus Routes. It is also used for in-house development integrated with ICOS for pre-disaster modules and mapping building-level data to activities and student impacts.
Data Storage & Integration	SQL Server / SQL Server Integration Services (SSIS)	Widely used for legacy ETL and data warehousing tasks.
	SLDS (Statewide Longitudinal Data System)	Past longitudinal data project (about 10 years ago) that failed due to lack of organizational buy-in and developer scheduling issues; remnants in data marts.

Category	Tool	Description
	Redgate SQL Monitor	Basic capabilities used for monitoring SQL Server performance; limited usefulness for analyzing performance.
	Access Databases	Still in use in certain departments for legacy reporting.
	Excel	Extremely prevalent across departments; used for data entry, validation, reporting, and reconciliation.
	PowerShell / Manual Scripting	Used in data transformations such as moving vendor Secure File Transfer Protocol (SFTP) files, and other file movement management.
	Comma Separated Values (CSV) files / Flat Files via SFTP	Common for cross-system data exchange between districts and OSPI.
	Application Programming Interface (API)	Employed for one program's nightly data handoff to external testing systems, facilitating demographic information outflow.
	C# / C++ Jobs	Nightly automated scripts that pull data from SFTP sites or send to external sites, separate from SQL Server execution.
	Stored Procedures	Dictate data processing and business rules in databases, updated annually for new job information and federal policy compliance.
	Washington Query / Query 2006	Core tool for ad-hoc data pulls and federal reporting (e.g., NCLB fact tables), with Statistical Product and Service Solutions (SPSS) syntax for downloads.
	SharePoint	Primary platform for sharing data files (e.g., CSVs from Student Information) and collaborative access across teams.
	System Drives (e.g., S-drive)	Legacy storage folders for file management.
Reporting & Visualization Tools	Tableau Server (not SaaS version)	Used widely for analytics, especially for public-facing analytics tools like Report Card.
	SSRS (SQL Server Reporting Services)	Used for standard reports and document storage.
	R	Programming language used by a few analysts for deeper data analysis; siloed to a few analysts and a couple of teams.

Category	Tool	Description
	Python	Utilized for data analysis within the ICOS console to evaluate efficacy, such as correlating Heating, Ventilation, and Air Conditioning (HVAC) scores with report card grades and enrollment.
	Power BI	Deployed in Power Platform for analytics and potentially small instances inside Fabric but not yet mature in self-service usage. Usage varies by department.
	Excel	Serves as a fallback visualization and reporting tool for users not running SQL scripts.
Student Education Systems	Comprehensive Education Data and Research System (CEDARS) / CEDARS Loader	OSPI's centralized education data repository system where districts submit student information; heavily customized.
	Individualized Education Program (IEP) Systems	Multiple solutions used across Local Education Agencies (LEAs) to submit special education data; not standardized statewide.
	School Apportionment System (SAFS)	Custom internal tools for financial data submissions and apportionment funding calculations; aging architecture with heavy manual processes.
Identity, Access, & Security	EDS	Main legacy custom-built authentication tool that hosts all OSPI's applications and central access for districts.
	Entra ID	Sparsely being used for applications outside of EDS; no current adoption strategy, Compliance requirements pending in 2025.
	CommVault	Used for daily backups on all servers; no data purging and data continues to accumulate year-after-year.
Productivity, Collaboration, and Project Tools	Microsoft Teams / Outlook / Office 365	Primary communication and productivity suite.
	Microsoft Forms	Used for surveys internally and externally for ad-hoc projects.
	Smartsheet	Used for automating in-house business processes such as contract management, project management and other data collection activities such as some transitioned iGrants activities that were not moved to Salesforce Grantmaking.
	Visio	Used for diagramming technical processes.

Category	Tool	Description
	Miro	Used for interagency collaboration "whiteboard" and project management visuals.
	Team Foundation Server (TFS)	Legacy tool being used for application development lifecycle, also used for IT ticketing, data pipeline tasks, and business rule implementation; currently planning a transition over to Azure DevOps.
	Azure DevOps	Database development teams have transitioned from TFS into Azure DevOps.
	File Maker	Personal desktop database tool being phased out used for internal business processes. Internal database for storing (Electronic Grants Management System) EGMS files, supporting compliance reviews and non-compliance summaries. Legacy tool that is phased out with adoption of EGMS.
	iGrants	Legacy grants management system replaced by EGMS; previously used for special education funding and reporting. Some historical data may show up from iGrants.
SaaS / Vendor Tools for Public-Facing Program Management	Salesforce Grantmaking	Currently called EGMS (Electronic Grants Management System) and replacing legacy iGrants system, used for federal reporting indicators, but clunky navigation and high learning curve, also lacking File Maker capabilities.
	Alchemer	Survey software for data collection (e.g., Capstone facilities reports, program grants); enables flexible question adjustments but lacks real-time validation.
	eVal	Vendor-managed tool for school districts teacher and principal evaluations.
	SalesForce	Used for a few departmental case management business processes.

Core Data Systems and Infrastructure

OSPI operates a complex data ecosystem that includes several primary data systems:

Comprehensive Education Data and Research System **(CEDARS)** functions as the central student information database collecting enrollment, demographic, and program participation data.

- Receives data from district SIS systems via file uploads with specific formatting requirements.

- Uses a "logical delete" validation process where entire submissions are rejected if error threshold exceeded.
- Batch processing cycles with significant dependency management opportunities

Educational Data Systems **(EDS)** Integrates with multiple downstream systems including assessment and reporting.

- Core authentication platform and interface for school districts.
- Houses 42+ applications and modules with varying security models.
- Single point of access for majority of district-facing applications.

School Apportionment and Financial Services (**SAFS**)

- Manages critical financial processes including district budgets (F-195), annual financial statements (F-196), initial budget planning (F-203), and personnel data reporting (S-275).
- Calculates and distributes state funding to school districts through the apportionment system.
- Rely on manual file uploads and batch processing for all financial data collection.
- Districts submit data through separate forms for each financial process.
- Manual validation is done regularly between related financial forms (F-195, F-196, F-203).

Educator Certification (eCert)

- Manages educator certification records and processes for approximately 75,000 active educators in Washington State.
- Processes initial certifications, renewals, endorsement additions, and interstate reciprocity for teachers, administrators, and educational staff associates.
- Maintains certification status data that determines educator eligibility to work in specific roles and subject areas within school districts.
- Exchanges data with the S-275 Personnel Reporting system and other OSPI educator databases through file-based transfers.
- Operates using file export and import processes for data exchange with district systems and other OSPI applications, without direct API connectivity.

Data Architecture & Integration

Current Architecture State:

- Heavily siloed systems with limited integration capabilities
- Manual data transfers common between systems
- Multiple separate SQL Server databases
- No centralized data warehouse architecture

- Many data marts not grounded in a governed Gold-layer breaks single-source-of-truth principle with multiple copies of "truth" that diverge over time
- Limited guardrails around ad-hoc tooling outside IT oversight create difficulties in governing data collection, processing, and analytics

Integration Methods:

- SFTP file transfers from districts and vendors
- Manual CSV exports/imports
- Manual surveys and forms through SaaS or EDS applications
- Nightly batch processing to copy data from web app servers to database servers for analytics

Data Validation:

- Logical Delete process in CEDARS rejects entire submissions for threshold violations.
- Manual verification steps between each system handoff
- Limited automated validation - primarily basic format checking
- Districts report resubmitting multiple times due to validation failures.

Technical Infrastructure

Database Systems:

- SQL Server on-premises (versions ranging from 2008 to 2022)
- Multiple separate databases per domain with no consistent naming conventions
- Limited cloud adoption
- Databases designed for thousands of records now have millions.

Development Environment:

- Legacy .NET Framework versions requiring updates
- Custom Windows Services without standardized logging or monitoring
- Outdated frameworks requiring extensive updates.
- Different technology stacks require specialized knowledge, creating key person dependencies.

Reporting Tools:

- Most visuals are on Tableau with minimal data modeling before being stored in Tableau Server
- Excel for data manipulation and completely disconnected from databases
- R and Python code for some analysis maintained by individual analysts
- Manual SQL queries with limited version control or documentation

- No consistent calculation definitions automation or enforcement across reporting tools

Additional Infrastructure Observations:

- Inconsistency in disaster recovery strategy across all critical systems and operations
- Authentication mechanisms vary by application and require standardization
- Network segmentation requires enhancement between environments
- Monitoring limited to basic server health - limited application performance monitoring
- Change management process exists but frequently bypassed for "emergency" fixes

People

OSPI's workforce is one of its greatest strengths, bringing deep institutional knowledge and a strong commitment to service. However, as the agency's technology environment evolves, staff capacity and skills must evolve with it. The current workforce model reflects years of experience maintaining legacy systems but has limited exposure to modern cloud platforms, data engineering, and automation practices. This section examines how existing roles, skills, and organizational structures support current operations and where growth training, new roles, and clearer governance will be needed to sustain modernization.

Workforce

The internal workforce at OSPI demonstrates commitment and institutional knowledge, but lacks exposure to many of the modern skills, roles, and tools needed to support cloud modernization at scale. There are a few key considerations that are widespread across state government:

- Limited staff time for training exists, as employees juggle full workloads, leaving little time to participate in training or skill-building.
- Role clarity remains limited, with overlapping responsibilities and unclear decision-making authority leading to some duplication of effort, gaps in data consistency, and operational inefficiencies.
- Budget constraints for hiring or training currently limit the ability to bring in specialized talent or expand training programs necessary to expand skillsets to support a cloud modernization at scale.
- Staff are eager for change and improvement across the organization, but multiple concurrent initiatives and evolving priorities due to budget and resourcing constraints have led to change fatigue and potential challenges to future state adoption.

Key considerations include:

- Skills coverage shows opportunities for expansion, as current staffing structures have not incorporated dedicated Cloud Architects, DevOps Engineers, Platform Engineers, Data Operations (DataOps) practitioners, or Financial Operations (FinOps) analysts
- The Platform Development persona show limited experience on average, with opportunities for skills building in areas such as infrastructure as code, containerization,

and automation frameworks that will support future modernization efforts and sustainment.

- CI/CD and automation capabilities are in early stages, with teams currently needing to develop experience in continuous CI/CD pipelines, IaC, and automation tools.
- While some exposure to cloud concepts exists, hands-on familiarity with Azure or other public cloud environments is in early stages. Most staff are grounded in on-premises operations.
- Professional development approaches require a more thorough approach, as the organization explores pathways to transition from current informal learning opportunities toward structured cloud enablement and development programs.
- Given the scope of cloud transformation, OSPI may need to augment with external partners in the near term to provide architectural, engineering, and operational support that the current team lacks resources for.
- Leadership engagement remains a key strength across OSPI, with opportunities for continued refinement of team structures and expansion of technical change management processes.

Gap Notes: OSPI will need a structured workforce development plan that defines required future-state roles, maps existing skills, and addresses gaps through training, hiring, or vendor support. Without this, modernization efforts will rely too heavily on overallocated resources, slowing near-term development and long-term sustainability.

Change Management

Change management practices are not fully documented and vary across teams, which creates risk to stability and consistency. Without this established practice, OSPI risks inconsistency in adoption and resistance management.

Key considerations include:

- **Approval processes** currently rely on collaborative approaches, with management reviews and decisions happening through direct conversations and email communications, while formal Change Advisory Board (CAB)structures are still a gap.
- **Tooling:** Systems such as ServiceNow or Jira are not used for ticketing and formalized workflows are not established.
- **Rollback Procedures:** Rollback and emergency change management processes are not documented.
- **Low Formalized Change Capability** with Business Process Model and Notation (BPMN) modeling (1.5), resistance management (2.5), and change frameworks (2.5) all underdeveloped.

Gap Notes: A formalized, standardized change control process with governance, workflows, and rollback procedures is not in place.

Communications

OSPI staff consistently report uncertainty about their roles in the modernization effort and how upcoming changes will impact day-to-day responsibilities. Organizational readiness remains limited by unclear communication channels and the absence of a structured change network. Without targeted communication and embedded champions, staff risk feeling disconnected from the modernization journey, which may slow adoption, reinforce resistance, or increase turnover.

Key considerations include:

- **Role Clarity:** Many staff do not understand how modernization affects their specific responsibilities, creating confusion and hesitation to engage.
- **Communication Channels:** Current updates are ad hoc and vary by team; no agency-wide cadence exists for progress reporting or feedback loops.
- **Message Framing:** Staff often perceive change as disruptive rather than supportive, with limited emphasis on "what's in it for me".
- **Change Network:** There is no structured group of sponsors, leads, and champions to localize messages, surface risks, or model adoption behaviors.

Gap Notes: Without structured communications and a change network, modernization will rely heavily on informal channels, leading to inconsistent understanding and uneven adoption across teams.

Process

OSPI's business processes form the foundation of how educational data moves through the organization, from collection and validation to reporting and analysis. While these workflows have enabled compliance and service delivery for years, they remain largely manual, fragmented, and dependent on individual expertise. The absence of consistent automation and governance limits efficiency and introduces risk in data quality and timeliness. This section examines these core operational processes and the governance structures that support them, leading into an assessment of OSPI's overall process maturity, data operations, and key obstacles that must be addressed to enable modernization.

Operations

OSPI's data modernization must support the full lifecycle of educational data across 100+ programs and processes, from initial collection through final reporting and analysis. These processes span multiple domains each with unique requirements but sharing common needs for accuracy, timeliness, and integration.

Data Collection and Validation

The foundation of OSPI's data ecosystem begins with collecting accurate information from 295 school districts through multiple channels and systems. CEDARS serves as the primary collection point for student-level data including enrollment, demographics, program participation, and course completion through batch submissions that require extensive validation and error

correction cycles. Personnel reporting through S-275 captures educator assignments, qualifications, and experience for over 60,000 certificated staff, linking teachers to students and courses while tracking compliance with certification requirements. Financial data collection via SAFS manages district financial reporting for over 15 billion in state funding, requiring complex reconciliation between student counts, staff allocations, and expenditure categories.

Program Eligibility and Compliance

Determining student eligibility and ensuring program compliance represents a critical process area that cuts across multiple systems and requires sophisticated business rules. Special Education determination involves referral tracking, evaluation management, IEP development, and service delivery monitoring for over 150,000 students with disabilities, with complex federal reporting requirements and strict compliance deadlines. Title I eligibility calculations require integration of poverty data, enrollment information, and school demographics to determine funding allocations and program participation, impacting both school-wide and targeted assistance programs. While English Learner identification and services track assessment results, program placement, and progress monitoring for students requiring language support, with annual reporting requirements and reclassification processes.

Funding and Apportionment

The calculation and distribution of education funding represent one of OSPI's most complex and critical process areas, directly impacting district operations and requiring absolute accuracy. Basic education funding calculations process enrollment data through complex formulas considering grade levels, vocational programs, and special populations to determine monthly apportionment payments to districts. Categorical program allocations layer additional funding streams including special education, transportation, learning assistance, and federal programs, each with specific eligibility rules and reporting requirements. Levy and Local Effort Assistance calculations integrate property valuations, voter-approved levies, and state equalization formulas to ensure funding equity across districts with varying local tax bases.

Assessment and Accountability

Managing statewide assessments and calculating accountability metrics requires coordinating with multiple vendors while maintaining strict security and validity standards. State assessment administration encompasses test registration, accommodation tracking, secure material distribution, and result processing for multiple grade levels and subject areas. Score processing and reporting involves complex matching algorithms to link assessment results to student demographics and program participation, enabling disaggregated analysis while protecting individual privacy. Accountability calculations transform raw assessment data into school and district report cards, computing complex metrics including growth models, graduation rates, and English learner progress while ensuring statistical validity and fair comparisons.

Reporting and Analytics

Transforming collected data into actionable insights for stakeholders represents the culmination of OSPI's data processes, requiring integration across all domains. Federal reporting obligations including EDFacts, Civil Rights Data Collection, and program-specific reports demand precise data integration and quality validation to avoid compliance penalties and ensure continued

funding. Legislative reporting responds to increasing demands for evidence-based policy analysis, requiring rapid turnaround on complex queries that span multiple years and data systems. Public transparency initiatives include school report cards, fiscal dashboards, and equity analyses that must present complex data in accessible formats while maintaining appropriate privacy protections.

Data Governance and Quality

Ensuring data accuracy, consistency, and appropriate use across all systems requires robust governance processes that currently rely heavily on manual effort. Data quality validation occurs at multiple points from initial district submission through final reporting, with each program maintaining separate validation rules and correction procedures that often conflict or duplicate effort. Privacy and security management must identify and protect personally identifiable information across all systems while enabling appropriate access for authorized users, a process complicated by varying sensitivity levels and access requirements across programs. Master data management for critical entities like students, educators, schools, and districts lacks centralized coordination, resulting in inconsistent identifiers and difficulty tracking individuals across programs and time.

Integration and Interoperability

Connecting disparate systems to enable comprehensive analysis and reporting represents an ongoing challenge requiring significant manual intervention. Cross-system matching links student records across enrollment, assessment, and program participation systems using probabilistic matching algorithms when unique identifiers are unavailable or inconsistent. Temporal alignment reconciles data collected on different schedules, daily attendance, monthly enrollment, annual assessments, to create point-in-time snapshots for analysis and reporting. External system integration manages data exchange with federal systems, other state agencies, higher education institutions, and third-party vendors, each with unique formats and requirements.

These processes, while currently fragmented across multiple systems and teams, must be reimagined as integrated workflows within the modernized architecture, enabling automation, ensuring consistency, and ultimately improving educational outcomes for Washington's students.

Standards

OSPI has varying standard documented procedures for how its technology systems are built and connected. Integration practices, architecture practices, and master data practices can vary based on which team did the work. Different teams follow their own rules for naming files, organizing data, and storing documents, which creates a lack of clarity and inconsistencies across the organization.

Key considerations include:

- Integration style relies heavily on batch processes and point-to-point connections, with limited evidence of standardized API-first or event-driven approaches.

- Data architecture fragmentation exists, as there is legacy enterprise data architecture in place. Systems operate independently, and no centralized or federated data access model exists to support cross-system reporting or analysis.
- Metadata and master data lack a unified metadata layer or master data strategy. This hinders data lineage, consistency, and consistent reporting across domains.
- Federation and cross-agency access show that current patterns do not support scalable data federation across internal or partner agencies, limiting the ability to share data securely and efficiently for statewide decision-making.

Gap Notes: Reference architectures and repeatable solution patterns are not currently defined or enforced.

Change Control

Current change control practices within the organization are inconsistently applied and lack a standardized framework, leading to heightened risk around system stability, auditability, and operational consistency. While individual teams have developed their own informal approaches, these methods vary significantly and are often undocumented, making it difficult to ensure accountability or repeatability. Without a unified governance model, changes may be implemented without sufficient oversight, introducing the potential for service disruptions, security vulnerabilities, or compliance gaps.

Key considerations include:

- **Approval Process:** Approvals may occur informally or through ad-hoc management review, conversations, and/or email, without a formal CAB.
- **Tooling:** Systems such as ServiceNow or Jira are not used for ticketing and formalized workflows are not established.
- **Rollback Procedures:** Rollback and emergency change management processes are not documented.

Gap Notes: A formalized, standardized change control process with governance, workflows, and rollback procedures is not in place.

Governance

Infrastructure Governance

Infrastructure governance structures are limited and inconsistently applied, resulting in fragmented practices across the organization. Policies and standards, such as naming conventions, tagging, cost management, and security, are not consistently defined or enforced, leaving critical areas open to interpretation. While IT governance appears centralized, the degree of departmental ownership is unclear, creating uncertainty in decision-making and accountability. Without formal enterprise-wide bodies, such as an Architecture Review Board or Data Governance Council, the organization faces risks of policy drift, inconsistent oversight, and reduced alignment with modernization and compliance goals.

Key considerations include:

- **Policies:** Naming conventions, tagging standards, cost management policies, and security standards are not consistently defined or enforced.
- **Organizational Model:** IT governance appears centralized, though the degree of federated ownership by departments is unclear.

Gap Notes: Enterprise-wide governance processes and bodies (e.g., Architecture Review Board, Data Governance Council) need to be formally established.

Data Governance

Data governance at OSPI lacks both automated systems and adequate staffing.

Key considerations include:

- **Governance Tools:** OSPI has no automated tools to manage data security, quality checks, or compliance monitoring. Governance activities must be done manually, creating significant workload and increasing risk of human error.
- **Policy Awareness:** While data governance policies exist in written form, most staff members are either unaware of these policies or do not consistently follow them in their daily work.
- **Roles and Accountability:** Documented data governance roles and responsibilities are lacking for data being democratized internally and externally, leading to unclear decision-making authority and data quality issues.

Gap Notes: OSPI's data governance framework lacks automated enforcement tools, which would enable organization-wide adoption of data management standards. Current governance efforts are manual, understaffed, and lack the authority needed for effective oversight.

Maturity Matrix

This Information Technology Infrastructure Library (ITIL) Maturity Matrix provides a framework for evaluating OSPI's current IT service management capabilities and defining the transformation journey toward optimized cloud-native operations. The five-level progression maps directly to the gaps and opportunities identified throughout the assessment.

1. Level 1 (Initial) is the foundational level where practices are disorganized and can be ad hoc or chaotic in nature.
2. Level 2 (Repeatable) represents implemented practices that are achievable using a basic set of defined activities that are already familiar to the team.
3. Level 3 (Defined) are when practices and patterns become well-organized, documented, and can be integrated with other practices in a systematic manner in alignment with the governance by design principle.
4. Level 4 (Quantitative) is the level of maturity that brings highly organized, measured, and monitored practices with regular reporting and optimization strategies built into the IT portfolio aligning with the self-service enablement principle.

5. Level 5 (Optimized) is the level that represents fully integrated and automated practices with continuous self-improvement capabilities that can run like clockwork aligning with the zero-touch operations principle.

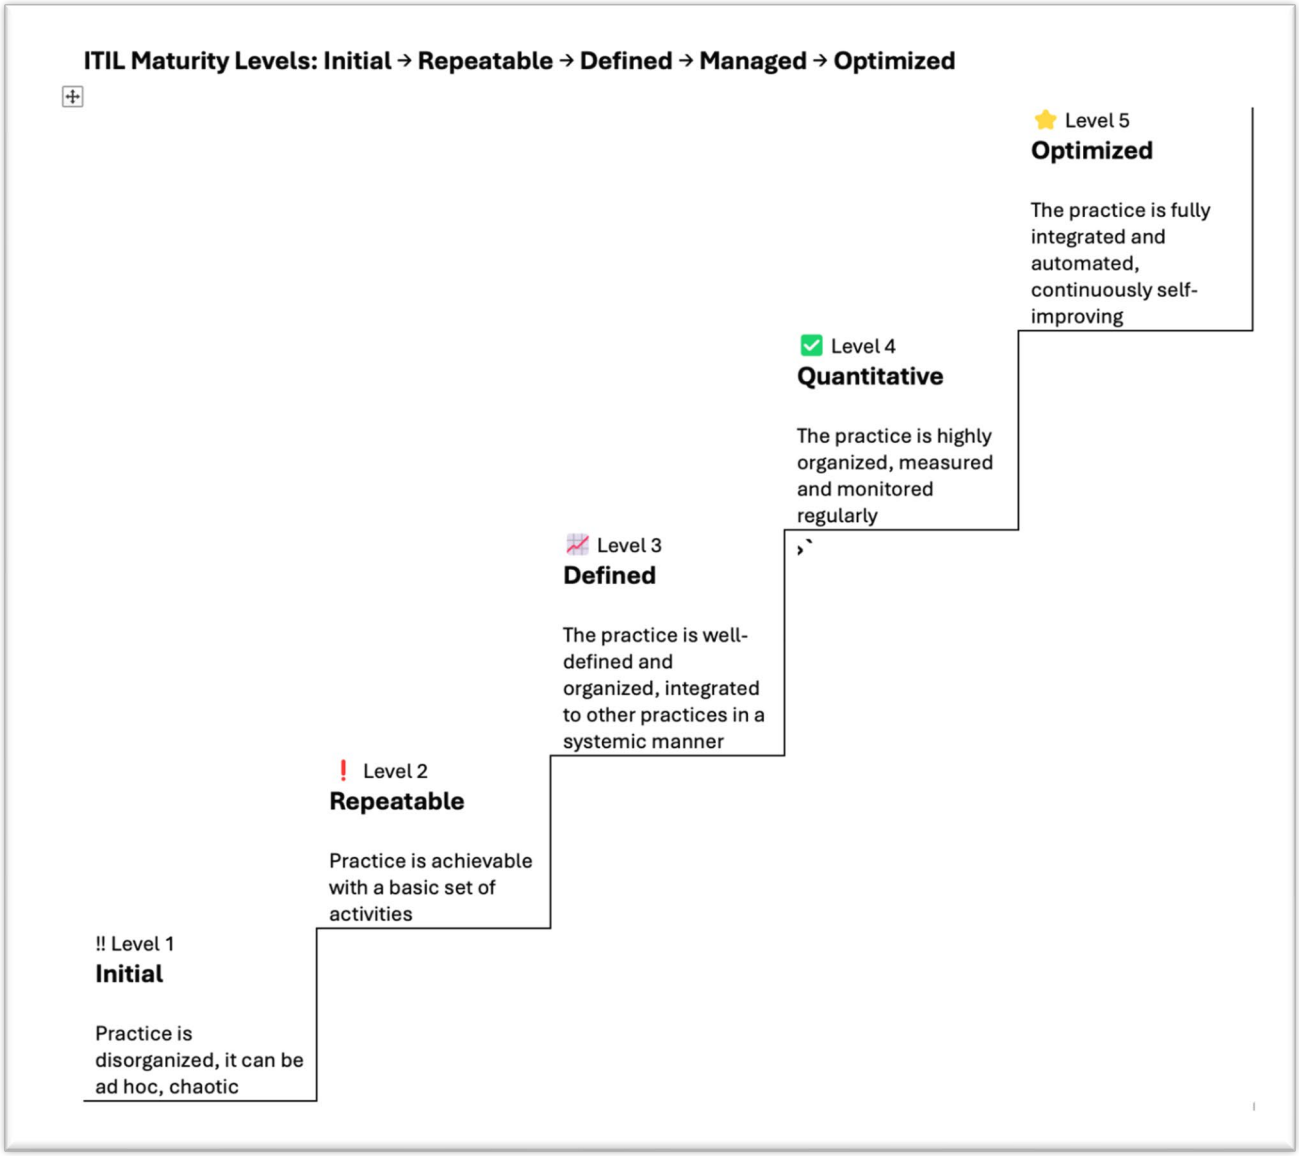

The following table provides a current state maturity assessment for OSPI across different architectural domains using the ITIL scale:

Architecture Domain	Current Maturity	Notes
Infrastructure	! Repeatable (L2)	Basic custom authentication + firewall, no automation
Repeatable (L2)	! Repeatable (L2)	Core apps known, dependencies unclear
Data Architecture	! Repeatable (L2)	Hundreds of ad-hoc databases, no centralized data lake, manual reconciliation

Architecture Domain	Current Maturity	Notes
Patterns	Initial (1)	Limited defined patterns, inconsistent usage
Processes	Initial (1)	Some processes defined, not standardized
Change Control	Initial (1)	Informal approvals, no CAB
Infrastructure Governance	Initial (1)	Policies not fully defined, ad-hoc practices
Data Governance	! Repeatable (L2)	Formal written governance, inconsistent enforcement

Agency-Specific Themes Scoring

The heatmap below is scoring OSPI's current state, from a scale of 1-5, using 23 strategic themes that apply across OSPI's infrastructure and data landscape. This same scoring will be used to evaluate the possible future state approaches and vendor solutions and will provide a direct comparison with OSPI's current state against future state solutions. The definitions of the 23 strategic themes can be found in Section 7.

Scoring (1-5):

1. **No Capability.** The capability does not exist in the current architecture, and OSPI cannot build a solution to meet the requirement with current staff.
2. **Buildable Partial Capability.** The capability does not exist in the current architecture, but OSPI can build a solution that partially solves the requirement (at no additional cost) with current staff.
3. **Buildable Full Capability.** The capability does not exist in the current architecture, but OSPI can build a solution that fully solves the requirement (at no additional cost) with current staff.
4. **Existing Partial Capability.** The capability exists in the current architecture but only partially solves the requirement (may require enhancements or customizations at no additional cost) with current staff.
5. **Existing Full Capability.** The capability exists in the current architecture and fully solves the requirement without the need for modifications.

Strategic Theme	Weight	Agency Current	Agency Weighted
Interoperability	0.15	1	1.15
Usability	0.10	2	2.10
Accessibility	0.15	1	1.15
Procurement Path	0.10	5	5.10

Strategic Theme	Weight	Agency Current	Agency Weighted
Data Visualization	0.10	3	3.10
Support & Maintenance	0.10	1	1.10
Manual Work Automations	0.15	2	2.15
Master Data Management	0.15	2	2.15
Managed Services	0.10	1	1.10
Compliance & Governance	0.10	2	2.10
Compute Fit	0.15	2	2.15
Database Compatibility	0.10	5	5.10
Storage Flexibility	0.10	1	1.10
High Availability	0.10	3	3.10
Disaster Recovery	0.10	2	2.10
Migration Tooling	0.10	2	2.10
Cost Optimization	0.10	5	5.10
Ecosystem & Support	0.10	1	1.10
Geographic Coverage	0.05	1	1.05
Innovation & Roadmap	0.05	2	2.05
Zero-touch Operations	0.15	1	1.15
Governance by Design	0.10	1	1.10
Self-Service Enablement	0.05	2	2.05
	SUMS	48.0	50.5

Data Operations & Process Challenges

Current Pain Points and Limitations

Initial feedback has shown challenges such as data silos, slow data request response times, and maintenance issues with aging infrastructure. While these systems have served OSPI reliably for decades, our analysis reveals that what appears as stability masks deeper structural challenges that limit OSPI's ability to meet modern educational data needs.

- **Critical Operational Impacts:** The current environment requires extensive manual reconciliation between systems, with staff spending considerable time validating data across multiple sources to ensure accuracy. What should be straightforward data requests often require days of investigation to find authoritative sources. Federal compliance reporting, while successfully completed, demands significant manual effort and coordination across teams. The absence of automated validation means quality issues are discovered late in processes, requiring reactive remediation rather than proactive prevention.
- **Systemic Technical Debt:** OSPI maintains numerous databases across multiple systems, each evolved to meet specific needs over time. While individually functional, these systems lack integration, resulting in the same data being calculated and stored differently across departments. Critical business logic exists in stored procedures and custom code written by staff no longer with the organization, creating knowledge gaps. The architecture, built incrementally over decades, makes comprehensive changes challenging without risking disruption to working processes.
- **Human Resource Constraints:** The technical complexity requires specialized knowledge that has accumulated in long-tenured staff members. This expertise, while valuable, creates vulnerabilities when key personnel are unavailable or retire. The absence of self-service capabilities means analytical staff serve as intermediaries for routine data requests, limiting their capacity for strategic analysis. Training new staff requires extensive mentorship periods due to undocumented business rules and system interdependencies.
- **Data Quality and Governance Gaps:** Without centralized governance, departments have developed their own definitions and validation rules, leading to inconsistent interpretations of the same metrics. The lack of transparent calculation documentation makes it difficult to explain how specific numbers are derived, particularly for complex federal formulas. When discrepancies arise between reports, finding the root cause requires manual investigation through multiple systems. Historical reproducibility is limited, making year-over-year comparisons and audit responses time-intensive.
- **District and School Impacts:** School districts navigate multiple submission systems with varying requirements and timelines. While these systems function, they require districts to maintain expertise in numerous platforms and data formats. Validation occurs post-submission, meaning districts may not discover issues until after critical deadlines. Smaller districts with limited technical staff face proportionally greater burdens in meeting complex submission requirements across multiple systems.

- **Strategic Limitations:** The current infrastructure, while operational, constrains OSPI's ability to provide prompt, comprehensive insights for policy decisions. Cross-domain analysis requires manual data compilation from multiple systems, limiting the ability to understand interconnected factors affecting student outcomes. The absence of predictive capabilities means interventions are reactive rather than preventive. Modern expectations for real-time data access and self-service analytics cannot be met within current architectural constraints.

Data Issue Analysis

During the Discovery Phase, over 800 issues were documented across OSPI collectively. Analysis shows that solving the first 6 issue categories would address over 77% of OSPI's challenges:

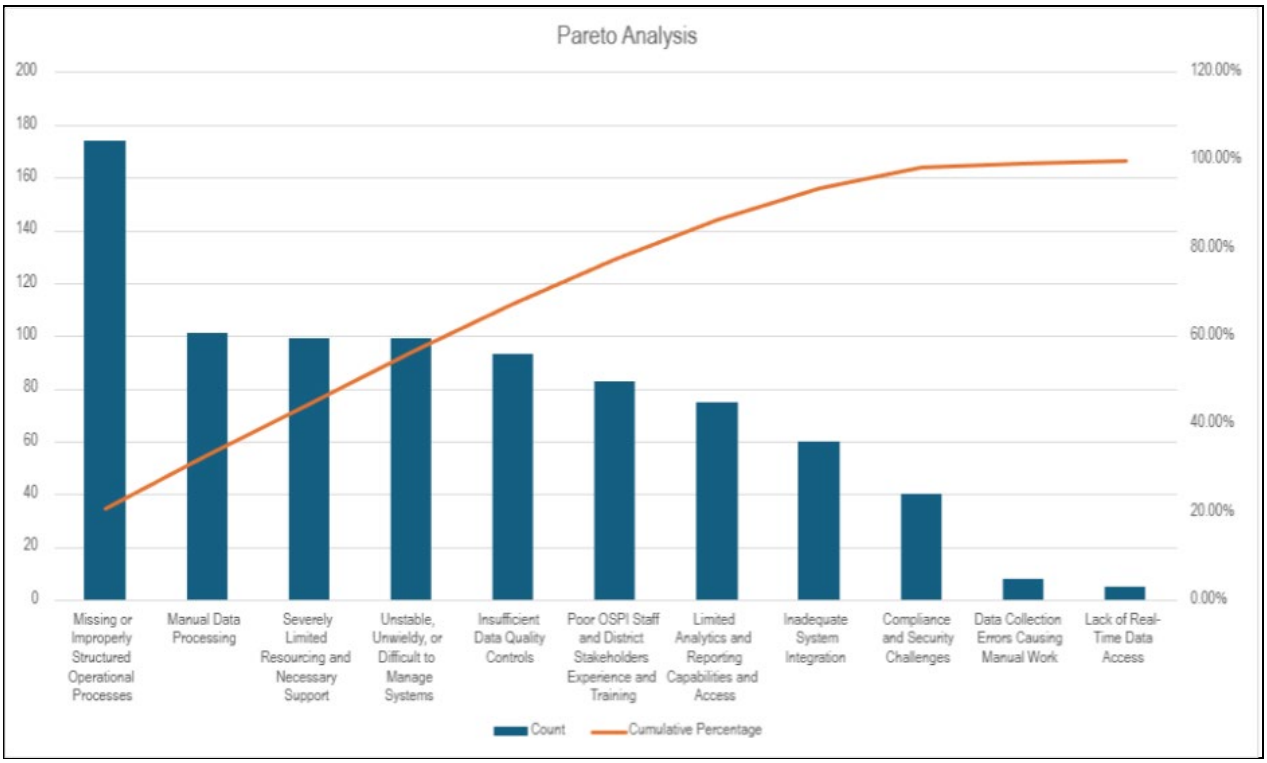

This analysis shows that if we solve the first 6 Issues Categories of:

- Missing or Improperly Structured Operational Processes
- Manual Data Processing
- Severely Limited Resourcing and Necessary Support
- Unstable, Unwieldy, or Difficult to Manage Systems
- Insufficient Data Quality Controls
- Poor User Experience and Training for OSPI Staff and District Stakeholders

It would solve over 77% of OSPI's issues. With the number one priority being "Operational Processes Missing or Improperly Structured."

The following are definitions of each issue category:

- **Missing or Improperly Structured Operational Processes -** Issues where business processes lack proper definition, automation, or governance, leading to inefficient operations, unclear responsibilities, or inconsistent execution. This includes missing workflows, undefined procedures, and poor change management.
- **Manual Data Processing -** Issues where data handling requires human intervention to move, transform, or process information between systems or stages, rather than automated workflows. This includes manual data entry, file transfers, calculations, and reconciliation tasks.
- **Severely Limited Resourcing and Necessary Support -** Issues where the resourcing is so inadequate that no matter how the processes and tools are improved, there are still going to be resourcing gaps that will still prevent the team from performing at optimal levels.
- **Unstable, Unwieldy, or Difficult to Manage Systems -** Issues related to system reliability, performance, user interface design, or maintenance complexity. This includes frequent outages, slow response times, outdated technology, and systems requiring excessive technical expertise to operate.
- **Insufficient Data Quality Controls -** Issues where systems lack automated validation, standardization, or error detection mechanisms, leading to inconsistent, incomplete, or inaccurate data. Normally leading to manual rework and multiple manual validation processes by multiple teams.
- **Poor OSPI Staff and District Stakeholders Experience and Training -** Issues where system interfaces are not intuitive, users lack necessary skills or knowledge, or adequate training resources are unavailable. This includes complex workflows, missing documentation, and insufficient support for users to effectively utilize systems.
- **Limited Analytics and Reporting Capabilities and Access -** Issues where users cannot easily analyze data, generate insights, or create reports due to system limitations, lack of tools, or restricted access. This includes inability to perform cross-domain analysis, generate custom reports, or access data for decision-making.
- **Lack of Real-Time Data Access -** Issues where data is not available immediately when needed, requiring users to wait for manual data pulls, batch processes, scheduled updates, or manual refreshes. This prevents timely decision-making and responsive operations.
- **Compliance and Security Challenges -** Issues related to meeting regulatory requirements (FERPA, CEDS, state laws), protecting sensitive information, maintaining audit trails, and ensuring proper access controls. This includes both technical security gaps and process compliance challenges.
- **Inadequate System Integration -** Issues where systems cannot share data seamlessly, requiring manual bridges, duplicate data entry, or preventing unified workflows. This

includes lack of APIs, incompatible data formats, and siloed applications that do not share automated data integration.

- **Data Collection Errors Causing Manual Work -** Issues where incorrect, incomplete, or improperly formatted data submissions create downstream work to identify, communicate, and correct errors. This is distinct from data quality controls as it focuses on the manual effort required to remediate collection mistakes.

Conclusion

OSPI's business environment shows that the agency's success depends on the alignment of its technology, people, and processes, each of which demonstrates both resilience and strain. The current technology landscape provides a relatively stable foundation but remains rooted in aging, on-premises systems that limit scalability, automation, and interoperability. OSPI's workforce brings deep expertise and commitment but lacks the modern cloud, data engineering, and automation skills necessary to sustain large-scale modernization. Meanwhile, core processes continue to rely on manual coordination and fragmented governance structures, slowing data flow and decision-making across programs.

Together, these conditions reveal an organization that has optimized for reliability but now faces diminishing returns without structural change. To meet the state's growing educational and technology expectations OSPI must modernize in a way that unites technology transformation, workforce development, and process redesign under a cohesive data strategy.

2.1.2 Business Needs

OSPI's current environment underscores an urgent need to evolve from maintaining legacy operations to delivering integrated, data-driven services that can adapt to Washington's dynamic education landscape. Modernization must address long-standing challenges in technology, workforce capability, and process efficiency to enable faster, more accurate, and equitable decision-making. By establishing a secure, unified data foundation and aligning people and processes around it, OSPI can shift from reactive compliance work to proactive insight generation and statewide collaboration.

The following business needs define the priorities that will guide this transformation and ensure that modernization delivers measurable value:

Enable Data-Driven Decision Making

There is an urgent need to transform OSPI from a reactive, compliance-focused organization to a proactive, insight-driven agency that can identify and address educational challenges before they become critical. A modern data platform will provide real-time visibility into enrollment trends, attendance patterns, achievement gaps, and resource allocation, enabling policymakers and educators to make timely, evidence-based decisions. This transformation requires breaking down data silos to create integrated views across student performance, educator effectiveness, and fiscal resources, allowing OSPI to identify successful interventions and scale them statewide.

Streamline Operational Efficiency

The current manual processes consume substantial staff resources that could be better directed toward strategic initiatives and direct support to districts. Modernized architecture will automate routine data collection, validation, and reporting workflows, reducing the time required for monthly apportionment calculations from weeks to days and federal reporting preparation from months to weeks. This efficiency gain extends to districts, which currently must enter the same data into multiple OSPI systems, creating redundancy and increasing the likelihood of errors that require time-consuming corrections.

Ensure Regulatory Compliance and Security

There is a critical need to address the growing complexity of state and federal compliance requirements while simultaneously protecting sensitive student and educator data from evolving cybersecurity threats. A modern data platform will embed compliance rules directly into data pipelines, automatically validate submissions against regulatory requirements, and maintain comprehensive audit trails for all data access and modifications. This approach transforms compliance from a reactive scramble during audit season to a continuous, automated process that reduces risk and ensures OSPI can confidently meet all regulatory obligations.

Improve Educational Equity and Outcomes

The inability to quickly analyze data across programs prevents OSPI from identifying and addressing equity gaps in real-time, potentially allowing disparities to persist for years before being discovered and addressed. A unified data architecture will enable sophisticated analysis of outcomes across demographic groups, geographic regions, and program participation, highlighting where additional support or intervention is needed. This capability is essential for OSPI to fulfill its commitment to equity and ensure all students, regardless of background or location, receive the support they need to succeed.

Enhance Stakeholder Experience and Transparency

School districts, educators, parents, and community members increasingly expect timely, accurate, and accessible information about educational programs and outcomes. A modern platform will provide self-service analytics capabilities, reducing the burden on OSPI staff while empowering stakeholders with direct access to relevant data. This transparency builds trust, enables local innovation, and supports collaborative problem-solving between OSPI, districts, and communities to improve student outcomes.

2.1.3. Business Opportunities

Modernization offers OSPI a path to move from maintaining disconnected legacy systems to leading a unified, secure, and data-driven education ecosystem. By aligning people, processes, and technology, the agency can strengthen equity, accountability, and operational efficiency across programs. Cloud-based automation, standardized APIs, and modern governance will reduce manual effort, improve data quality, and deliver real-time insights that inform funding, compliance, and policy decisions.

These opportunities create both organizational and technical capacity for OSPI to operate as a connected enterprise which empowers staff, simplifies district interactions, and enhances transparency for students, families, and lawmakers.

Together, they establish the foundation for evolving existing roles and creating new growth opportunities that build a future-ready agency.

Persona Transformation Journey Through Modernization

- Program Manager Future Journey
    - Instead of logging into multiple systems, program managers will access integrated dashboards showing real-time program performance, automated compliance tracking, and predictive insights about program effectiveness. Their journey transforms from data gathering to strategic analysis and proactive intervention.
- Data Analyst Future Journey
    - Analysts will spend 80% of their time on actual analysis rather than data preparation, using automated pipelines, pre-built data models, and self-service tools that enable rapid hypothesis testing and insight generation. Their journey shifts from repetitive report building to innovative problem-solving.
- District Administrator Future Journey
    - Superintendents, data administrators, and principals will receive timely, integrated reports that combine academic, financial, and operational data in clear, actionable formats. They will access improved public dashboards showing school and district performance metrics and trend analyses prepared by OSPI's data teams. Their journey evolves from piecing together fragmented reports to receiving comprehensive insights that support evidence-based planning and resource allocation decisions.
- Parent/Student Future Journey
    - Families will benefit from enhanced public reporting that provides clearer, easily understood, more comprehensive information about school and district performance, program offerings, and educational outcomes. School report cards will be updated more frequently with easier-to-understand visualizations, potentially available in multiple languages, and include contextual information that helps families make informed decisions. Their journey transforms from searching across multiple websites and outdated PDFs to accessing current, integrated information about their schools and districts through improved public dashboards and reports. Missing Future State Personas

The current organizational structure lacks several critical personas necessary for successful data modernization:

- **Architecture Design Team** will be providing long-term planning for system integration, establishing enterprise standards and ensuring new investments build toward a coherent future state rather than creating new silos.
- **Platform Development Team** will build and maintain the modern infrastructure, implementing infrastructure as code, automation, and integrated security that current IT staff lack capacity to deliver.
- **Data Development Team** will focus specifically on data platform engineering, building robust pipeline automation, implementing data quality and data security frameworks, and creating reusable data products, machine-learning products, and analytics assets that democratize insights across the organization for internal and external stakeholders, especially to support existing analyst teams.
- **Organizational Change Team** are essential for driving adoption, developing training programs, and redesigning business processes to leverage new capabilities rather than perpetuating old workflows in new systems.

2.1.4 Business Service Goals

The business service goals for OSPI modernization focus on strengthening the agency's internal operations while improving how services extend to school districts, Educational Service Districts (ESDs), and educational partners. The primary emphasis is on building consistent, automated, and secure processes that reduce manual effort, improve data quality, and ensure reliable delivery of core services.

The primary service goals:

- Reduce redundant data handling across systems by automating manual processes and standardized workflows.
- Ensure trusted and transparent reporting by establishing unified data governance, automated validation and interoperability between core systems.
- Improve funding calculations, compliance and accountability with modernized apportionment and financial service systems.
- Reduce outages and protect sensitive information with cloud-based, scalable, and compliant platforms that create a stronger infrastructure.
- Provide districts and ESDs with streamlined tools, intuitive interfaces, and real-time access to information.
- Support sustainable operations, future state competencies, and professional pathways by equipping OSPI staff with modern tools, training, and defined roles.
- Generate timely and accurate insights that inform data-driven decisions for OSPI leadership and policy makers.

2.1.5 Statutory Requirements

OSPI is guided by a set of federal and state statutory requirements that define how student data, financial resources, and educational services must be managed. Federal mandates establish strict expectations for individualized education planning, accountability frameworks, data privacy, and accessibility. Washington state law further codifies OSPI's responsibilities in apportionment, special education, staffing, and reporting. These legal obligations establish compliance thresholds that directly influence the design and operations of OSPI's technology systems. As OSPI modernizes its IT environment, statutory requirements serve as foundational guardrails to ensure that all technology decisions uphold equity, accountability, and transparency for Washington's students, families, and educators.

Requirement	Description
44 U.S.C. 3551	IT systems must align with NIST, continuous monitoring, and incident response.
WaTech 141.10, 141.20	Mandate NIST aligned IT security, modernization must integrate zero-trust, encryption, and vulnerability management.
RCW 42.56.590, 19.255.010	Data breach notification laws require monitoring, logging, and incident response in IT systems.
44 U.S.C. CH. 29, 31, 33	Requires proper record retention/disposition. Systems must include archiving, defensible deletion, and audit trails.
RCW 43.105, 43.105.210	Accessibility and governance for IT systems must comply with WCAG 2.0 AA.
20 U.S.C. 1232g	Protects student privacy by requiring strict identity management, encryption, and audit logging.
WAC 392	Defines reporting for enrollment, personnel, certifications where IT systems must standardize and automate submissions.
20 U.S.C. 6301	Mandates student performance and equity reporting which drives the need for interoperable data systems and dashboards.
RCW 28A.505	Requires transparent financial/apportionment reporting, IT systems must handle funding with accuracy and auditability.
RCW 28A.300	Defines requirements for data collection, reporting, and accountability which require secure and scalable IT systems.
RCW 42.56	Washington State Public Records Act requires IT systems to have efficient retrieval, redaction, and disclosure management.
20 U.S.C. 1400	Ensures students with disabilities receive appropriate services requiring IT systems to support IEP creation, tracking, and reporting.

Requirement	Description
RCW 28A.155	Implements IDEA in Washington requiring systems to document and monitor IEPs and engagement.
42 U.S.C. 2000d	IT must support multilingual access and equity monitoring.

SECTION 3 – OBJECTIVES

Modernization represents an opportunity to align agency operations, technical growth, and statewide educational outcomes under a secure unified data-driven vision. The business value extends beyond IT upgrades to deliver strategic people-centered benefits that improve efficiency, equity, and accountability across the educational system.

To realize these outcomes, business objectives should focus on enhancing service delivery, improving security and regulatory compliance, while building the organizational and technical capacity needed to sustain modernization over time. There are several key objectives to support the initiative:

- **Streamline service delivery** by simplifying access to agency programs and data through interoperable cloud-based scalable systems and consistent user experiences.
- **Automate compliance and reporting** by implementing automated validation, audit trails, and workflow-based monitoring that ensures timely and accurate submissions.
- **Enhance data quality and integration** with an interoperable architecture that enables real time insights and cross agency analytics.
- **Increase operational efficiency** through modern security-first infrastructure and the adoption of automation to reduce maintenance overhead, improve reliability, and free staff for higher value work.
- **Advance security, privacy, and accessibility** with embedded zero-trust principles, privacy by design, and WCAG/ADA compliance across all platforms ensuring equitable and secure service delivery
- **Support data-driven decision making** by providing stakeholders with actionable, high-quality information to inform funding and accountability measures.
- **Strengthen governance and sustainability** through data governance frameworks that standardize processes, promote transparency, and guide ongoing system improvement.
- **Workforce enablement and readiness** through integrated organizational change management to build staff capacity, ease adoption of new systems and processes, and empower continuous improvement.

3.1.1 Service Delivery Enhancements

Service delivery modernization transforms how OSPI fulfills its mission to provide equitable, data-driven services statewide. Shifting from a fragmented legacy operations to an integrated, efficient, and user-centered service model that builds upon its business service goals and objectives. Together they improve efficiency, transparency, and user experience while ensuring equitable access to accurate and secure information statewide.

Service delivery impacts through modernization:

Enhanced Oversight and Compliance

- Automate compliance monitoring across federal and state requirements through governance by design
- Provide real-time audit trails and lineage for funding, certification, and student data which reduce risk for penalties.

Data-Driven Policy and Resource Allocation

- Shift from fragmented reports to unified, real-time dashboards that directly inform legislative and funding decisions.
- Enable predictive analytics to anticipate enrollment trends, educator shortages, or program performance gaps.
- Improve accuracy and timeliness of apportionment and other fiscal distributions.

Improved Governance, Communication, and Collaboration

- Establish a cross-functional Data Governance Council that gives OSPI a stronger role as a statewide convener.
- Create shared reference architectures, data definitions, and collaboration platforms across divisions and with ESDs.
- Standardize decision-making through Architecture Decision Records (ADRs), reducing silos and shadow IT.

Workforce Modernization and Capacity Building

- Upskill OSPI technical staff in cloud, DevOps, DataOps, and AI/ML integration.
- Build institutional resilience by reducing single points of failure from "institutional knowledge" systems.
- Provide staff with tools to spend more time on analysis and support, less on manual data work.

Operational Efficiency and Risk Reduction

- Automate manual data transfers, validation, and reconciliations.
- Establish disaster recovery, continuity, and zero-trust security foundations.
- Reduce maintenance of legacy systems and technical debt.

Transparency and Accountability

- Provide legislators, auditors, and the public with consistent, reproducible data backed by automated lineage.
- Ensure internal OSPI decisions (budget allocations, program monitoring, staffing) are transparent, and data driven.

- Align modernization directly with OSPI's four strategic goals.

Strategic Innovation and Agility

- Position OSPI as a national leader in education data modernization by adopting scalable, AI-enabled data architecture.
- Enable rapid response to legislative data requests.
- Build a culture of continuous improvement with iterative enterprise architecture practices.

3.1.2 Response to Statutory Requirements

OSPI's efforts to modernize provide an opportunity to meet and exceed statutory obligations by embedding compliance, accessibility, and data integrity within every system and process. Federal and state laws governing education, data privacy, and accessibility require reliable, timely, and auditable systems. This initiative fulfills these mandates by automating compliance functions, standardizing data management, and providing transparent reporting that supports equitable education outcomes which improves service delivery, reduces administrative burden, and builds stronger public trust. Several key alignment themes include:

- **Compliance by Design** including automated validation, lineage tracking, and audit trails ensure adherence to federal and state reporting requirements.
- **Accessibility and Equity** to ensure all systems provide equitable access for staff, educators, families, and students.
- **Transparency and Accountability** supporting real-time reporting obligations for open data and efficient government.
- **Data Security and Privacy** based on zero-trust architecture and automated governance for data protection.
- **Operational Efficiency** through modern platforms to reduce manual effort and risk in alignment with the goals of efficient state technology investments.
- **Cross-Agency Integration** with improved interoperability supports data-sharing mandates with Education Research and Data Center (ERDC) and other state agencies while maintaining compliance.

The objectives and service delivery enhancements position OSPI to operate as a unified, responsive, and secure digital organization. By aligning modernization with statutory mandates and strategic goals, OSPI strengthens its capacity to deliver equitable, transparent, and efficient data-driven services.

This strategic alignment will create lasting organizational capacity ensuring technology investments continually improve educational outcomes and operational performance. These advancements lay the groundwork for sustained impact across the agency and the broader education system.

SECTION 4 – IMPACTS

The impacts of modernization will be felt across multiple layers of the OSPI ecosystem. It will strengthen collaboration and alignment with partners through consistent data standards and reporting frameworks. Internal capacity will increase/grow by improving governance, enhancing staff skills, and streamlining core internal operations. Greater efficiency and reduction in duplicative manual work across all programs will enable a more responsive service delivery and deliver better outcomes for students statewide.

The diagram below simply shows how the future state data warehouse can directly support internal users (inter-agency) and external users (intra-agency) from easily accessible curated data. Internal users can interact with Gold-level data directly from interfaces such as Power BI or Azure App Services and pass on this curated data to external users; additionally, external users can leverage Gold-level data surfaced in public-facing applications such as Tableau embedded in Azure App Services.

Inter-agency and Intra-agency support quick view diagram:

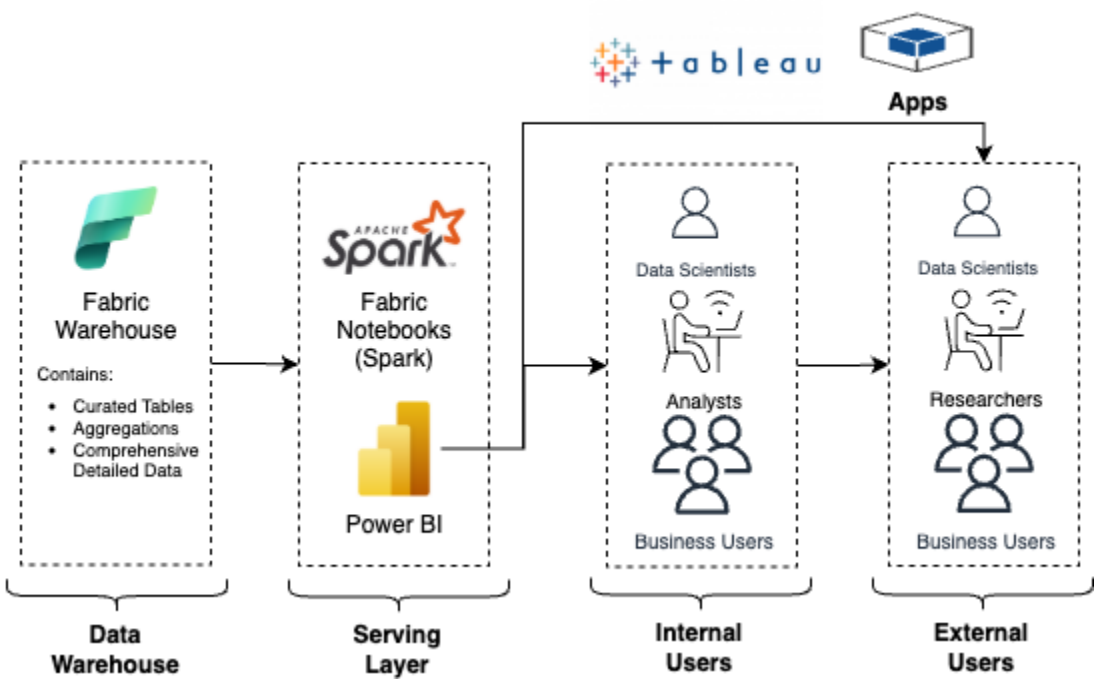

4.1.1 Inter-agency (within OSPI)

OSPI's internal divisions often work in isolation, with different teams managing data separately and making decisions without a full picture. The impacts of this fragmentation show up in reporting delays, duplicated work, and misaligned strategies. Key issues include:

- Siloed data ownership prevents collaboration across divisions.

    - Program areas like Student Services, Assessment, and Finance each manage data independently, with limited visibility into one another's systems or reporting workflows.
- Duplicated reporting effort wastes staff time and introduces inconsistencies.
    - Multiple teams request the same data from Local Education Agencies (LEAs) or maintain redundant spreadsheets, leading to conflicting versions and confusion about which data is authoritative.
- Strategic planning suffers from disconnected insights.
    - Without a shared platform or coordinated governance, OSPI divisions make decisions based on partial views of the data landscape, limiting their ability to align initiatives or respond to statewide trends.

4.1.2. Intra-agency (external partners)

OSPI's external partners depend on timely and accurate data to make decisions. But gaps in access, visibility, and documentation create real challenges. These include:

- External stakeholders experience inconsistent access to OSPI data.
    - Districts, policy analysts, legislators, and the public rely on different entry points to OSPI data, each with varying levels of timeliness, usability, and completeness.
- Manual report requests delay decision-making.
    - Partners often wait weeks for OSPI teams to fulfill ad hoc data requests that could be self-served if the right tools and access controls were in place.
- Lack of transparency weakens trust in state reporting.
    - Without clear data lineage or documentation, LEAs may question how data is transformed between submission and reporting, which reduces confidence in statewide analyses and funding decisions.

4.1.3 Programs (public-facing services)

OSPI's public-facing programs stand to experience transformative improvements in data quality, accessibility, and reporting through the modernized platform, directly impacting educational outcomes and resource allocation across Washington State.

Assessment and Accountability

The modernized platform could eliminate manual data reconciliation between testing vendors, student information systems, and reporting platforms. Automated demographic crosswalks for subgroup reporting would ensure that achievement gaps are identified accurately and quickly, allowing districts to deploy targeted interventions before students fall further behind. Historical assessment data would become immediately accessible for longitudinal analysis, enabling

educators to track individual student growth over years rather than viewing single point-in-time snapshots, ultimately improving instructional decisions and student outcomes.

Special Education Services

The platform could consolidate fragmented IEP data currently scattered across district systems into a unified view of services, accommodations, and progress monitoring, ensuring that the state's special education students receive consistent services regardless of which district they attend. Automated data validation would ensure federal reporting accuracy while reducing the staff hours each district spends annually on manual reports, allowing special education coordinators to focus on student services rather than paperwork. Service providers would gain the ability to track outcomes across districts when students move, eliminating any unnecessary gaps in service continuity that affects mobile students and potentially preventing regression in skills.

Student Enrollment and Records

The architecture could create a single authoritative source for student enrollment, eliminating discrepancies in different enrollment numbers for the same schools in different records that currently creates cascading analytical issues for internal and external users needing accurate enrollment data. Automated deduplication could identify students enrolled in multiple districts simultaneously that can cause inaccurate duplicate funding claims while ensuring accurate per-pupil funding allocations that directly impact classroom resources.

Financial Reporting and Compliance

Federal program reporting could become automated through integrated data pipelines that would pull from authoritative sources rather than manual spreadsheet compilation. Title I allocations, E-rate calculations, and grant reporting could draw from the same validated datasets, potentially eliminating conflicting numbers across reports and reducing audit findings. This would save staff time, improve consistency across submissions, and reduce the risk of compliance issues that could affect funding or trigger corrective actions.

Multilingual Education Data

Transitional Bilingual Instruction Program and Title III reporting could shift from manual aggregation of district spreadsheets to automated data integration, reducing reporting timelines enabling program coordinators to identify and respond to emerging needs among the state's English learners before academic gaps widen. The architecture could enable tracking of students across districts to measure long-term English proficiency outcomes, providing the longitudinal data necessary to evaluate program effectiveness and adjust instructional approaches that could accelerate language acquisition timelines.

Nutrition Program Eligibility

The platform could automate the matching of enrollment data with income eligibility information, eliminating any unnecessary manual verification processes that could delay benefits of eligible students. Real-time data updates would ensure students maintain benefits immediately when moving between schools, preventing delays in providing necessary food

programs to students. Aggregate reporting could become instantaneous, enabling nutritional programs to adjust food orders and staffing faster to reduce annual food waste while ensuring no child goes hungry.

Attendance and Truancy Tracking

Real-time attendance data could flow directly from school systems into statewide dashboards, replacing monthly batch uploads that could delay intervention opportunities, allowing schools to implement support services while students are still engaged rather than after chronic absenteeism patterns are established. Architecture could enable pattern detection across schools to identify systemic issues such as transportation problems or health outbreaks, enabling district and state leaders to deploy resources proactively rather than reacting after attendance rates have already impacted learning outcomes and funding.

Graduation and Transcript Data

The platform could consolidate course completion, credit accumulation, and graduation requirement tracking across all districts into a unified repository, eliminating unnecessary manual transcript evaluation processes that could delay course enrollment for transfer students and can result in misplaced students repeating courses or missing graduation requirements. Real-time graduation cohort tracking could replace legacy annual calculation processes, enabling counselors to identify at-risk students months earlier and implement interventions that could improve the state's graduation rate. This unified system would ensure that the students who transfer between Washington districts maintain accurate progress toward graduation.

4.1.4 Subprograms (supporting services)

Internal support programs could experience significant operational improvements through data modernization:

Grants and Financial Management

Automated grant distribution could become possible through real-time enrollment and program participation data, potentially eliminating delays and errors in funding allocations. The architecture may enable complete transparency in tracking state funding from allocation through expenditure, with automated compliance reporting capabilities and comprehensive audit trails that could satisfy federal and state requirements. Manual reconciliation of grant expenditures against program participation could be replaced with automated validation.

Certification and Credentialing

The educator certification system could potentially integrate with district HR systems to automatically verify credentials and track professional development requirements. The architecture may enable automated alerts for expiring certifications, reducing the risk of non-compliance. Data modernization could eliminate duplicate data entry between OSPI certification systems and district personnel records, while providing real-time visibility into certification status across the state.

Transportation Services

Student transportation data could become integrated with enrollment and attendance systems, potentially enabling dynamic route adjustments based on actual ridership patterns. The platform may provide analytics on transportation costs per student, route efficiency, and service gaps. Automated data flows could eliminate manual reconciliation between transportation claims and student enrollment, potentially reducing funding errors.

Data Reporting and Compliance

The architecture could consolidate hundreds of separate district data submissions into unified, automated data flows. Federal reporting requirements (EDFacts, IDEA, Title programs) could pull from the same validated datasets, potentially eliminating conflicting statistics across reports. Districts may be able to submit data once for multiple compliance purposes rather than preparing separate submissions for each program.

Professional Development Tracking

The platform could potentially create a unified repository of educator training records, clock hours, and professional development activities. This may enable automatic verification of continuing education requirements and could provide analytics on training effectiveness correlated with student outcomes. Manual tracking of workshop attendance and credit hours could be replaced with automated systems.

Enrollment Projections and Planning

The architecture could enable sophisticated predictive analytics for enrollment forecasting by integrating demographic data, housing permits, and historical enrollment patterns. This may improve budget planning, staffing projections, and facility needs assessment. Real-time enrollment data could replace annual October counts for more responsive resource allocation.

Audit and Monitoring Functions

The platform could provide automated audit trails for all data changes, potentially reducing the effort required for compliance reviews. Risk-based monitoring could become possible through analytics that identify unusual patterns or outliers in program data. Manual sampling for audit purposes may be replaced with comprehensive automated reviews of all transactions.

Inter-district Transfers and Choice Programs

The architecture could potentially streamline the tracking of students participating in school choice, Running Start, and inter-district transfer programs. Automated data sharing between districts may eliminate manual paperwork for transfer students, while ensuring accurate funding follows the student. The platform could provide real-time visibility into capacity and enrollment for choice programs.

4.1.5. Customers of Agency Activities (end users)

Modernization delivers tangible benefits to every stakeholder in Washington's education ecosystem:

Students and Families

Parents gain unified access to all their children's educational information through a single secure portal, including grades, attendance, assessments, special services, meal accounts, and transportation. Real-time notifications keep families informed of academic progress, attendance issues, or required actions. Students among the one million served by OSPI benefit from more coordinated support services, faster intervention when struggling, and seamless transitions between schools or programs.

Educators and Administrators

Teachers access comprehensive student profiles showing academic history, intervention effectiveness, family engagement, and support services, eliminating the need to search multiple systems or request records. Automated administrative tasks free educators to focus on instruction, while predictive analytics help identify effective teaching strategies for specific student populations.

Researchers and Policy Makers

ERDC and approved researchers gain access to de-identified, longitudinal datasets linking early learning through workforce outcomes. The platform enables sophisticated research into educational equity, program effectiveness, and long-term student outcomes while maintaining strict privacy protections. Office of Financial Management can track education funding effectiveness and make data-driven budget recommendations.

Community Organizations

Local nonprofits, after-school programs, and community service providers can securely access relevant student information (with appropriate permissions) to coordinate wraparound services. This creates a true "village" approach to supporting student success, with all stakeholders working from the same accurate, timely information.

Tribal Communities

The six state-tribal education compact schools gain culturally responsive data management capabilities that respect tribal sovereignty while enabling participation in state programs. The platform supports collection of culturally relevant metrics and outcomes important to tribal communities while maintaining compliance with state and federal requirements.

The cumulative effect of these impacts creates a virtuous cycle where better data leads to better decisions, which improve outcomes, generating more valuable data for continuous improvement. This transformation positions Washington State as a national leader in educational data management, setting new standards for how states can leverage technology to improve educational outcomes for all students.

SECTION 5 – ORGANIZATIONAL EFFECTS

The shift to a modern data platform will have substantial effects on OSPI's workforce, operations, and delivery structure. These changes, while implemented through technology, affect how work is organized, how teams collaborate, and how services are delivered across the agency. The following subsections outline the expected organizational impacts and key considerations.

5.1.1. Impact on Work Processes

Workflows across the agency will shift from siloed reporting and manual data handling to coordinated, repeatable processes built on shared infrastructure. OSPI staff will interact with governed data products, standardized pipelines, and self-service tools rather than relying on ad hoc queries or isolated reports.

Routine data validation, reconciliation, and extraction tasks will become automated or eliminated. In their place, teams will spend more time analyzing trends, informing decisions, and responding to stakeholder needs. Over time, this will reduce data duplication efforts and improve OSPI's responsiveness to internal and external data requests.

5.1.2. Change Management and Training Needs

The scale of this change will require intentional support from the OSPI workforce. OSPI will need to implement a change management approach that includes clear communication, training, and support during and after implementation.

Staff across the agency will need training aligned to their roles, technical teams will need upskilling in cloud platforms, data governance, and DevOps, while analysts and business users will require orientation to the new data tools and access models. Program staff will need to understand how to request data, submit feedback, and participate in governance processes.

This training effort should be staged to support onboarding, day-to-day operations, and long-term adoption. Equally important is providing accessible documentation, channels for feedback, and mechanisms to monitor where additional support may be needed.

5.1.3 Job Content

As legacy systems are phased out and manual processes are streamlined, staff responsibilities will shift. Data analysts will focus less on sourcing and reconciling data and more on generating insights. Developers will move from maintaining point-to-point integrations to building scalable, governed pipelines.

New roles will emerge, such as Data Product Managers, Platform Engineers, and Governance Stewards, while other roles will expand in scope. These shifts will increase the strategic value of many positions and require clear definitions of responsibilities, expectations, and success measures.

In many cases, existing staff will be able to transition into new roles with appropriate training and support. A small number of specialized skills may need to be sourced externally or filled through targeted hiring.

5.1.4. Impact on Organizational Structure

The future-state operating model restructures OSPI's delivery approach around eight data product teams, each accountable for a specific domain (e.g., Student, Finance, Assessment). These teams will be cross-functional and operate with end-to-end responsibility for the delivery, maintenance, and governance of their assigned data products.

This model replaces the current structure of functionally siloed teams and shadow IT processes with integrated teams aligned to strategic priorities. Each product team will manage its own backlog, iterate its releases and embed governance as part of its daily operations.

To support this shift, OSPI will need to stand up shared roles and establish governance structures that provide coordination without centralizing control. As these teams mature, the model will allow OSPI to deliver data faster, with greater transparency and accountability across the organization.

SECTION 6 – PROPOSED SOLUTION

ISG determines that it is feasible for OSPI to implement cloud solutions from any of the major providers: Azure, Amazon Web Services, or Google Cloud. The key differentiator is OSPI requires a modern data strategy to manage 100+ terabytes of enterprise data currently residing in on-premises SQL Server systems and decades of tech debt. The existing infrastructure lacks automated governance, quality validation, and self-service capabilities, creating operational bottlenecks and critical advanced analytics limitations. The core future state recommendation that would align most with OSPI's needs would be a data architecture that can support three key architectural principles: Zero-Touch Operations, Governance by Design, and Self-Service Enablement. The proposed solution below outlines a release-by-release agile product approach that balances immediate operational needs with long-term scalability requirements.

Before diving into the journey from current state to future state, as mentioned, below are details about the key architectural principles that make the future state recommendation, address the needs of OSPI, and solve the issues found in section 2. These key principles are the driving pillars of the future state data architecture leveraging the medallion process as the core storage and data quality framework.

This section is organized into three core areas that together define the path forward: Technology, which outlines the architecture and tooling; People, which addresses skills, training, and adoption; and Process, which defines the modernization methodology and phased implementation strategy.

6.1.1. Technology Recommendations

The future-state architecture addresses OSPI's critical infrastructure and data platform needs through a cloud-agnostic, scalable design anchored in automation, governance, and accessibility. The following section explains the key architectural principles and technical foundation that underpin the proposed solution.

Key Architectural Principles

The three key architectural principles form the foundation of a truly modern, enterprise-grade data platform that transcends traditional IT limitations and meet the needs of OSPI's current legacy problem and help build a future-proof solution.

By implementing Zero-Touch Operations (or come as close as possible), OSPI will eliminate the operational burden that typically consumes 70-80% of data team resources, freeing them to focus on innovation rather than maintenance.

Governance by Design ensures that security, quality, and compliance aren't afterthoughts or checkboxes, but are seamlessly woven into every data interaction making governance an enabler rather than a barrier to data accessibility.

Self-Service Enablement democratizes data access across the organization, transforming data from a technical asset managed by IT into a business asset leveraged by everyone.

Together, these principles create a virtuous cycle:

Key Data Architectural Principles Virtuous Cycle Diagram

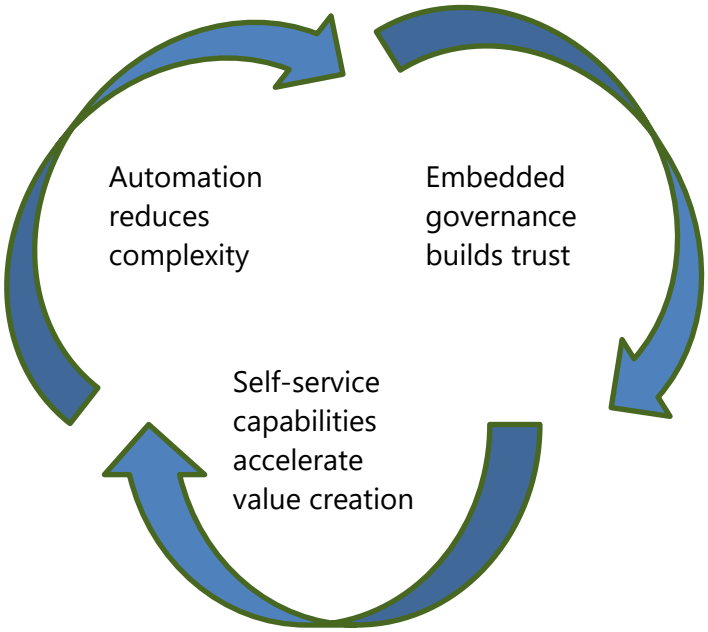

This approach fundamentally shifts the data platform from a cost center requiring constant manual intervention to a strategic accelerator that scales with the business, maintains compliance automatically, and empowers users at every level to make data-driven decisions in real-time. The result is a platform that is simultaneously more powerful and simpler to use while delivering enterprise capabilities.

1. Zero-Touch Operations

Zero-touch operations directly solve OSPI's "Missing or Improperly Structured Operational Processes" and "Manual Data Processing" issues that were found as the first and second most prevalent issue categories in the Pareto Analysis, along with the third and fourth categories. This principle represents autonomous infrastructure that eliminates manual intervention through intelligent automation across all platform components.

The capability encompasses event-driven orchestration where pipelines activate instantly upon data arrival, utilizing file watchers, message queues, and change data capture triggers to eliminate scheduling dependencies and ensure real-time processing. Security and governance controls implement policy inheritance and propagation mechanisms that cascade automatically through data lineage, enabling role-based access and encryption policies to flow seamlessly from source systems through to consumption layers without manual configuration. The platform enforces embedded quality standards where data quality rules, schema validation, and business logic apply transparently at the point of ingestion without requiring explicit configuration steps or manual intervention from technical teams. Infrastructure resources benefit from dynamic optimization that scales elastically based on real-time workload patterns, automatically

provisioning compute capacity during peak demand periods and deallocating resources during idle times to ensure cost efficiency.

2. Governance by Design

Governance by design directly solves OSPI's "Insufficient Data Quality Standards" and "Compliance and Security Challenges" issues that were found as the fifth and ninth most prevalent issue categories in the Pareto Analysis. This principle establishes comprehensive data stewardship embedded into every layer of the architecture, ensuring compliance and control without impeding operational efficiency.

The architecture implements automated metadata capture where every dataset is instantly cataloged with technical, business, and operational metadata, creating a searchable knowledge base that serves as the single source of truth for organizational data assets. Automation models power intelligent data classification that automatically identifies and tags sensitive information including personally identifiable information, protected health information, and financial data to ensure appropriate handling and compliance with regulatory requirements.

The platform maintains end-to-end lineage tracking that provides complete data journey visualization from source systems through all transformations to final reports, enabling comprehensive impact analysis and rapid root cause investigation when issues arise. Continuous quality monitoring operates through real-time data profiling with anomaly detection, completeness checks, and business rule validation that triggers proactive alerts before data quality issues can propagate through downstream systems.

The entire framework operates on a zero-trust security model implementing policy-driven access controls with encryption at rest and in transit, comprehensive audit logging, and automated compliance reporting built directly into the data flow.

3. Self-Service Enablement

Self-service enablement directly solves OSPI's "Poor OSPI Staff and District Stakeholders Experience" and "Limited Analytics and Reporting...Access" issues that were found as the sixth and seventh most prevalent issue categories in the Pareto Analysis along with the rest of the prevalent issue categories. This principle democratizes data access by empowering business users with intuitive tools while maintaining robust governance controls throughout the organization.

The capability provides frictionless data discovery where business users can browse, search, and request access to datasets through intuitive data catalogs without submitting IT tickets or requiring technical knowledge of underlying systems. Natural language processing powers conversational analytics that enable users to ask questions in plain English and receive immediate visualizations, insights, and recommendations without understanding query languages or technical specifications.

The platform delivers automated insight generation through report builders reducing time-to-insight from weeks to hours and enabling rapid decision-making. User-friendly interfaces can be built from Gold-layer data that enable program managers to build reporting solutions independently without requiring programming expertise.

A semantic business layer provides pre-built metrics, calculated fields, and standardized business definitions that ensure consistent interpretation and calculation methodologies across all consumer touchpoints and reporting interfaces.

What Is This Architecture?

This cloud-agnostic modern data platform architecture provides a proven blueprint for building enterprise-scale data solutions across any cloud provider. It codifies industry best practices into a repeatable framework that maximizes automation and governance while eliminating operational complexity and manual overhead.

Core Design Philosophy

"Automation First" drives every architectural decision which creates self-operating, self-healing, and self-governing systems. This approach eliminates traditional bottlenecks where manual data management, security configuration, and governance enforcement create delays and require extensive teams.

The Seven-Layer Architecture Stack

1. Foundational Infrastructure (Cloud Landing Zone)

- Secure network foundation with isolated VPCs/vNets and private endpoints
- Centralized identity management through SSO and role-based access control
- Enterprise encryption with managed key vaults and automatic rotation
- Unified observability through centralized logging, monitoring, and compliance scanning
- Outcome: Establishes secure, governed cloud foundation that meets enterprise standards

2. Governance Control Plane (Central Hub)

- Policy automation engine enforcing rules across all data assets
- Continuous compliance validation with automated remediation
- Centralized audit aggregation for forensics and regulatory reporting
- Intelligent alerting with context-aware incident management
- Outcome: Maintains control and visibility without manual intervention

3. Ingestion Layer

- Multi-pattern data capture supporting batch, streaming, CDC, APIs, and file drops
- Inline validation with automatic quality checks and schema enforcement
- Smart error handling using quarantine zones and retry logic
- Sensitive data isolation with automatic PII/PHI detection and masking
- Outcome: Reliably captures data from any source without manual configuration

4. Storage Layer (Medallion Architecture)

- Bronze Zone: Immutable raw data preserving original fidelity
- Silver Zone: Cleansed, validated, and standardized datasets
- Gold Zone: Business-ready aggregations and feature stores
- Specialized Stores: Purpose-built for time-series, graph, document, and cache patterns
- Outcome: Progressive data refinement with clear quality levels

5. Processing Layer

- Elastic compute engines for distributed batch and stream processing
- Serverless functions triggering on events without infrastructure management
- Integrated ML/AI platforms for model training and inference
- Intelligent orchestration with dependency management and automatic retries
- Embedded quality gates ensuring data integrity at every transformation
- Outcome: Scalable processing that adapts to workload demands

6. Serving Layer

- High-performance query engines with OLAP acceleration
- API gateway exposing REST and GraphQL endpoints
- Semantic business layer translating technical data into business concepts
- Self-service interfaces through BI tools, notebooks, and natural language
- Outcome: Democratized access meeting diverse consumption patterns

7. Governance Components

- Automated data catalog with business and technical metadata
- Visual lineage tracking from source to consumption
- Continuous quality monitoring with drift detection and scoring
- Dynamic access controls with attribute-based policies
- FinOps integration for cost attribution and optimization
- Outcome: Trust through transparency, compliance through automation

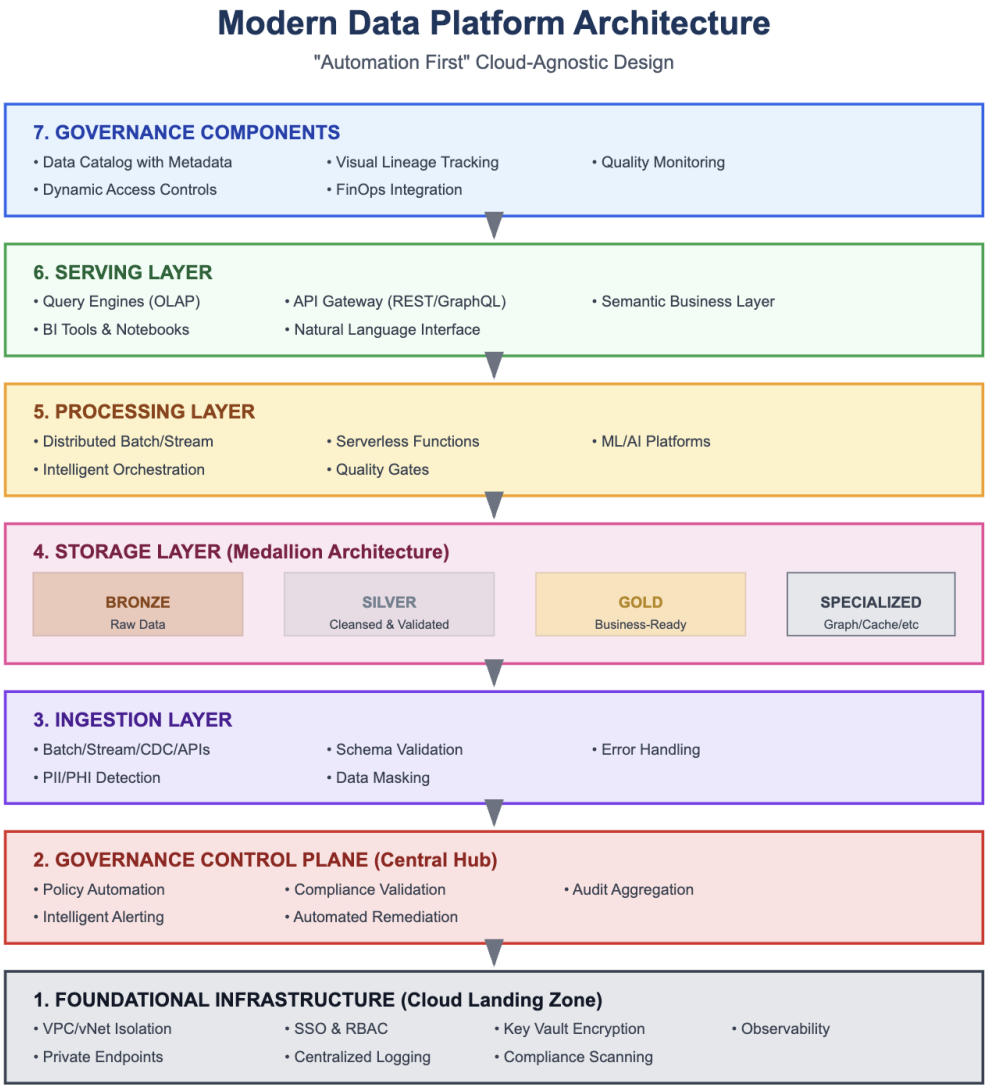

Future State Tools

Below is one set of tools recommended for OSPI's future state leveraging an Azure stack. Following this same blueprint, OSPI can choose to have a similar stack with any other cloud provider.

Category	Tool	Description
Cloud Platform Foundation	Microsoft Fabric F64+ capacity / Azure Databricks auto-scaling clusters	Elastic computing resources that scale up during peak periods and down during quiet times.
	Terraform / Bicep	Infrastructure-as-code enabling consistent, repeatable deployments.

Category	Tool	Description
	Azure Kubernetes Service (AKS) / Azure Container Instances	Container-based architectures for portable, efficient application deployment.
	Azure Functions / Azure Event Grid / Fabric Reflex	Event-driven serverless functions that execute only when needed.
Cloud Data Storage	Azure Data Lake Storage Gen2	Centralized data lake storage supporting any data type or format.
	Delta Lake / Azure Synapse Analytics	Transactional storage layers providing database-like consistency on files.
	Azure Synapse Dedicated SQL Pools / Microsoft Fabric Warehouse	Cloud-native data warehousing for high-performance analytics.
	Azure Cosmos DB / Microsoft Fabric KQL Database	Real-time operational data and analytics.
Data Orchestration and Integration	Azure Data Factory / Synapse Pipelines / Fabric Data Pipelines	Workflow orchestration engines managing complex dependencies.
	Azure Event Hubs / Azure Stream Analytics / Fabric Eventstream	Real-time streaming platforms for immediate data availability.
	Azure API Management / Microsoft Fabric API for GraphQL	Centralized API management ensuring granular data processing and control.
	Azure Data Factory Change Data Capture / Fabric Mirroring	Change detection capabilities capturing updates as they occur
Development and Operations	Azure DevOps Repos / GitHub	Distributed version control enabling collaboration.
	Azure DevOps Pipelines / GitHub Actions	Continuous integration and deployment pipelines.
	Azure Test Plans / Azure Load Testing	Automated quality assurance and testing.
	Azure Monitor / Application Insights / Log Analytics	Comprehensive system monitoring and alerting.

Category	Tool	Description
	Microsoft Fabric / Azure Machine Learning	Leveraging platforms for data product lifecycle management.
Data Governance and Security	Microsoft Purview / Unity Catalog / Fabric Data Governance	Unified governance platforms managing policies across all data.
	Microsoft Purview Data Map / Information Protection	Automated discovery and classification of information, especially isolation of sensitive data.
	Azure Policy / Microsoft Defender for Cloud	Dynamic security controls adjusting to policies and threats.
	Azure Privileged Identity Management (PIM) / Azure RBAC	Least privileged access management with time-bound permissions.
	Azure Monitor Activity Logs / Microsoft Purview Audit Logs	Complete audit trails tracking all data access and changes.
Analytics and Business Intelligence	Power BI	Self-service visualization platforms empowering end users.
	Azure Synapse Spark Pools / Azure HDInsight / Fabric Spark	Distributed processing frameworks handling large-scale computations.
	Azure Machine Learning with R/Python / Microsoft R Server	Statistical programming language integration for advanced analysis.
	Power BI Embedded / Azure Synapse Link	Analytics embedded directly into operational applications.

Below is a depiction of how the recommended future state tools would flow together in one architecture.

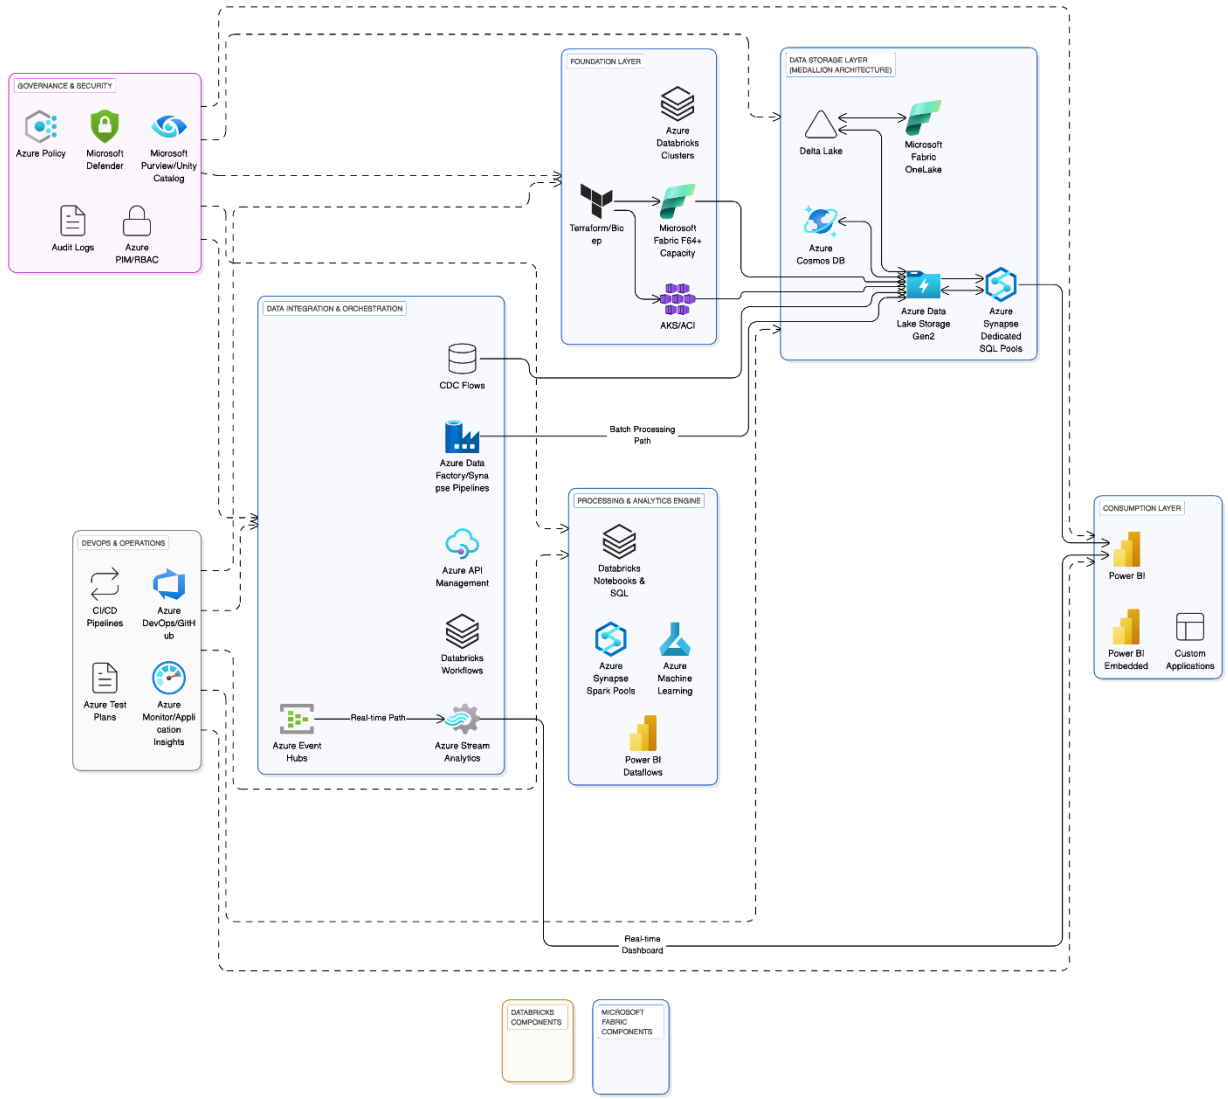

Future State Tools Diagram

This comprehensive Azure and Microsoft toolset provides OSPI with a good foundation for transitioning from legacy, siloed data systems to a modern, cloud-native architecture. The platform addresses critical modernization challenges through elastic scalability (auto-scaling Fabric capacities and Databricks clusters), unified data storage (ADLS Gen2 and OneLake eliminating data silos), and real-time capabilities (Event Hubs and Stream Analytics enabling immediate insights). Organizations can choose between two powerful paradigms, Microsoft Fabric for an all-in-one SaaS experience or Azure Databricks for a more customizable data lake approach, while maintaining enterprise-grade governance through Purview and Unity Catalog. The infrastructure-as-code approach (Terraform/Bicep) ensures repeatable, consistent deployments, while integrated DevOps tools (Azure DevOps/GitHub Actions) enable continuous delivery and monitoring. Most importantly, this ecosystem provides seamless integration between all components, reducing complexity and technical debt while enabling both technical teams (through notebooks and SQL) and business users (through Power BI) to extract value from data. This architectural approach transforms data from a cost center into a strategic asset by

reducing operational overhead through serverless and managed services, enabling self-service analytics, and providing flexibility to handle everything from batch processing to real-time streaming in a single, governed platform.

While technology defines the architecture, it is OSPI's people who will operate, govern, and make decisions using this platform. The future-state recommendations outlined below depend on aligning the right skills, roles, and supports to ensure successful adoption.

6.1.2. People Recommendations

To accelerate data modernization while maintaining agility and compliance, the initiative prioritizes the deployment of Microsoft Fabric in an F64+ configuration, supporting from 10TB to 30TB of data processing. This capacity facilitates immediate operations, allowing for seamless ingestion and analytics without the overhead of multiple SKUs. Complementing this, all processing logic is recommended to be developed using portable PySpark code, ensuring vendor-agnostic portability for future migrations – critical in dynamic environments where platform shifts may arise due to evolving needs or costs. Migration thresholds should be set to activate at 30TB data volume or when batch jobs exceed four hours towards Databricks infrastructure to support scaling performance.

Governance is to be centralized via Microsoft Purview, with automated data discovery, classification, and security on ingestion. This enforces metadata standards, lineage tracking, and access controls uniformly. To support parallel run during the migration, ExpressRoute will provide secure, low-latency connectivity from on-premises SQL servers, mitigating risks of data transfer delays and ensuring hybrid resilience during the highly complex transition from on-premises to cloud.

Analysts will be able to query governed and standardized datasets without second-guessing or wrangling data for weeks. Real-time governance dashboards in Purview will enable proactive compliance, monitoring, flagging anomalies contradicting compliance rules (e.g., FERPA, HIPPA) and will instantly alert data stewards and data owners.

The goals of the OSPI future state data platform are:

1. To scale effortlessly at 10X volume
2. Achieve near zero unnecessary operational manual touches
3. Apply governance immediately at source ingestion and throughout the data flow

The objective is to deliver sustained ROI through reduced maintenance and faster decision-making.

For OSPI staff and leadership to fully adopt and operationalize the platform, targeted investments in training, organizational change, and vendor partnerships will be essential.

Incorporated Skill Development ensures successful adoption

- Build out capability mapping for each OSPI user persona to identify what each role needs to do their job effectively. A capability map helps visualize those needs across the organization so gaps can be prioritized and mitigated.

- Create personalized learning paths.
- Provide hands-on training and practice.
- Hire specialized trainers where needed.

Vendor Selection considers OSPI goals alignment with a strategic evaluation

- Consider vendors with implementation plans that deliver speed-to-value.
- Prioritize tools that automate most manual processes, with top preference for automating security.
- Consider vendors with roadmap for frontend, data, and infrastructure.

6.1.3. Process Recommendations

Technology and people must be integrated through a repeatable, scalable process that ensures successful migration and continuous value delivery. OSPI's modernization strategy relies on a phased, agile implementation model with full value achieved in each release.

Modernization Approach

ISG recommends an agile modernization strategy built on rapid learnings through end-to-end releases that provides standalone value with each release. For each release, the team would have the opportunity to learn a little more of each of the following principles:

- Governance and Strategy Foundations – Establish data governance policies, roles, and metadata cataloging; build stakeholder alignment.
- Data Platform Architecture & Cloud Migration – Stand up cloud landing zones, develop a scalable data lakehouse architecture, and migrate priority datasets.
- Analytics Enablement & Advanced Insights – Deploy BI and self-service tools, implement KPI-driven dashboards, and integrate advanced analytics capabilities.
- Optimization & Continuous Improvement – Enhance automation, monitor data quality, and enable ongoing AI/ML innovation.

A phased migration approach minimizes disruption

- Start with small, high-value data projects using modern tools. Treat each step in the established data modernization roadmap (as described in Section 10 under Data Migration) – such as migrating a core system or enabling self-service for a program area – as a standalone release that delivers end-to-end value without requiring full system replacement.
- With each release, develop reusable components for the Feature Store – a centralized library of cleaned and engineered data features - that can be templated and used repeatedly.
- Gradually improve application of modern tools on previously deployed projects.
- Maintain parallel operations during transition.

- Provide ongoing and incorporated training and support.
- Each project should achieve full end-to-end modernization to build momentum.

6.1.4. Recommendations Summary

The proposed solution delivers an actionable blueprint for OSPI to transition from fragmented, manual, and compliance-constrained data operations to a modern platform that is automated, governed, and accessible. The future state addresses legacy technical debt through cloud-native architecture, reduces risk through embedded governance, and enables users with self-service analytics. Aligning the right technology, building internal capabilities, and executing data modernization through value-driven agile implementation will allow OSPI to modernize its data systems in a way that is scalable and sustainable over time.

SECTION 7 – ARCHITECTURAL APPROACH AND VENDOR ASSESSMENT

This section assesses possible implementation approaches and evaluates potential vendors to support OSPI's data modernization effort. It provides a high-level framework for phased execution, identifies key activities and milestones, and includes a comparative assessment of cloud vendors and implementation partners. The goal is to provide OSPI with a practical lens for planning, decision-making, and procurement preparation ahead of RFP development.

The approach analysis portion describes characteristics, advantages, and trade-offs between a home-grown, standalone architecture, and the adoption of a partner vendor which brings architectural capabilities as part of their core offering.

To support the vendor assessment, the feasibility study defined a set of strategic evaluation criteria aligned to OSPI's priorities, modernization goals, and known constraints. These criteria helped frame tradeoffs between platform options and clarify what capabilities will be most critical for successful implementation. Scores are summarized in each respective vendor section.

7.1.1. Approach Analysis

OSPI faces a critical decision in selecting the optimal implementation approach for their data platform modernization. Two distinct strategies offer different paths to achieving the same transformative outcomes, each with unique advantages and considerations that must be weighed against organizational capabilities, timeline constraints, and long-term strategic objectives.

Approach 1: Standalone Model

The Standalone Model represents a ground-up construction of the entire data platform, where OSPI builds each component from foundation to finish. Analogous to assembling a custom home from raw materials, this approach provides complete control over every architectural decision, technology selection, and implementation detail. The development would leverage OSPI's internal technical teams, potentially augmented by external vendor resources, to create a bespoke platform perfectly tailored to Washington's unique educational data ecosystem.

Key Characteristics:

- Custom-built architecture designed specifically for OSPI's requirements from scratch
- Direct integration with existing state systems and infrastructure
- Full ownership of intellectual property and source code
- Complete flexibility in feature prioritization and development sequencing
- Deep institutional knowledge embedded in platform design

Advantages:

- Total Customization: Every feature, workflow, and interface can be designed to match OSPI's exact specifications and unique requirements
- No Vendor Lock-in: Complete independence from external vendors for critical infrastructure
- Unlimited Scalability: Architecture can be designed to accommodate any future growth or requirement changes
- Intellectual Property Ownership: OSPI owns all code, documentation, and processes developed
- Deep Integration Capability: Seamless connection with legacy systems and state-specific requirements
- Institutional Knowledge Building: Internal teams develop comprehensive expertise in all platform components

Considerations:

- Extended Timeline: Full platform development typically requires 3-5 years for initial implementation
- Higher Initial Investment: Significant upfront costs for development, infrastructure, and team building
- Resource Intensity: Requires substantial internal technical resources or expensive external consultants
- Ongoing Maintenance Burden: OSPI assumes full responsibility for updates, security patches, and feature development
- Implementation Risk: Higher probability of delays, scope creep, and budget overruns
- Talent Dependency: Success depends on recruiting and retaining specialized technical talent

Approach 2: Vendor Platform Model

The Vendor Platform Model leverages pre-built Data Platform-as-a-Service (DPaaS) solutions that provide immediate functionality out of the box. Comparable to configuring pre-built furniture using a kit that requires assembly but comes with all necessary components, this approach utilizes established platforms like the Onix Wingspan Bird Suite running on Google Cloud or Cloudwick Amorphic Platform running on Amazon Web Services. These enterprise-grade solutions have been tested across multiple implementations and can be rapidly configured to meet OSPI's needs.

Key Characteristics:

- Pre-built, proven platform architecture with established best practices
- Rapid deployment with configuration rather than development

- Vendor-managed infrastructure and platform maintenance
- Built-in compliance and security certifications
- Regular feature updates and enhancements included

Advantages:

- Accelerated Time-to-Value: Initial platform operational within weeks, full implementation in months
- Reduced Implementation Risk: Proven solutions with established deployment methodologies
- Lower Total Cost of Ownership: Shared infrastructure costs and included maintenance reduce long-term expenses
- Immediate Best Practices: Benefit from industry-standard approaches refined across multiple deployments
- Continuous Innovation: Regular platform updates provide new capabilities without additional development
- Reduced Technical Burden: Vendor handles infrastructure, security updates, and platform maintenance

Considerations:

- Customization Constraints: Some OSPI-specific requirements may need workarounds or process adjustments
- Vendor Dependency: Reliance on vendor for critical infrastructure and ongoing support
- Licensing Costs: Ongoing subscription fees that may increase over time
- Data Sovereignty Concerns: Need to ensure compliance with state data residency requirements
- Integration Complexity: May require adapters or middleware for legacy system connections
- Limited Control: Platform roadmap and feature prioritization influenced by vendor and other customers

Decision Framework

The choice between approaches should consider several key factors:

- **Timeline Urgency:** If OSPI needs immediate improvements to address critical compliance or operational issues, the Vendor Platform Model provides faster relief. The Standalone Model is better suited when there's runway for deliberate, phased development.

- **Budget Profile:** The Standalone Model requires higher upfront capital investment but may offer lower long-term operational costs. The Vendor Platform Model spreads costs over time through subscriptions but provides more predictable budgeting.
- **Technical Capabilities:** Organizations with strong internal technical teams and architecture expertise may prefer the control of the Standalone Model. Those seeking to focus on outcomes rather than technology management benefit from the Vendor Platform Model.
- **Risk Tolerance:** The Standalone Model carries higher implementation risk but lower long-term vendor risk. The Vendor Platform Model reduces implementation risk but introduces vendor dependency considerations.
- **Innovation Needs:** If OSPI anticipates pioneering new educational data practices unique to Washington, the Standalone Model provides necessary flexibility. If OSPI is looking to rapidly innovate and get off the ground running, the Vendor Platform Model delivers proven solutions faster.

Recommendation Considerations

Both approaches can successfully deliver the modernized data platform OSPI requires. The Standalone Model offers maximum control and customization at the cost of time, resources, and risk. The Vendor Platform Model provides rapid implementation and proven reliability while accepting some constraints on customization and control. The decision ultimately depends on OSPI's specific priorities regarding timeline, budget allocation, risk tolerance, and long-term strategic vision for their data ecosystem.

These approaches lead directly into vendor-specific analysis, as each provider offers unique functional and technical capabilities, pricing models, architectural approaches, and varying levels of platform support or enterprise flexibility.

7.1.2. Vendor Analysis

OSPI evaluated three options for modernizing its data architecture: AWS with Cloudwick's Amorphic platform, GCP with Onix's Wingspan platform, and building a custom solution on Microsoft Azure. The analysis reveals a critical decision point between adopting a modern data platform versus attempting to build comparable capabilities internally.

The evaluation assessed 23 strategic themes, listed below, weighted by importance to OSPI's transformation goals, with particular emphasis on automation capabilities, governance, and self-service enablement. Both platform solutions (Amorphic and Wingspan) demonstrate comprehensive existing capabilities that can be deployed within months, while the Microsoft native build approach would require years of development.

Strategic Theme Definitions

Area	Strategic Theme	Definitions
User Experience & Usability	Usability	The degree to which a system can be used by specified users to achieve goals effectively, efficiently, and with satisfaction.
	Accessibility	The design of a system or architecture so that all levels of technical expertise can use it, including non-technical team members.
	Self-Service Enablement	Allowing users to perform tasks (e.g. analyze data) independently without waiting on a technical team or needing IT intervention.
	Data Visualization	The graphical representation of information and data using visual elements.
Integration, Data Quality, & Governance	Interoperability	The ability of different information systems, devices, and applications to access, exchange, integrate and cooperatively use data.
	Master Data Management	The process of defining and maintaining consistent definitions of business entities.
	Compliance & Governance	Ease of adherence to regulations and internal policies for data and system management.
	Governance by Design	Building governance requirements into system architecture from start to finish meeting or exceeding security and compliance requirements.
Infrastructure & Architecture Fit	Compute Fit	The alignment of computing resources with workload requirements.
	Database Compatibility	The ability of different database systems to work together effectively.
	Storage Flexibility	The ability to scale and adapt storage solutions as needed.
	High Availability	System design ensuring operational continuity with minimal downtime.
	Disaster Recovery	Ease of automated procedures for recovering from system failures or catastrophes.
	Migration Tooling	The robustness of tools and processes for moving data between systems, especially from on-prem to cloud with parallel run.

Area	Strategic Theme	Definitions
Operations & Automation	Zero-touch Operations	Fully automated processes requiring no human intervention.
	Manual Work Automations	The conversion of manual processes into automated workflows.
	Support & Maintenance	Ongoing technical assistance and system updates to ensure continued operation.
	Managed Services	Third-party operation and management of IT processes, especially accompanying a tool that was procured.
Cost & Procurement	Cost Optimization	Strategic management of IT expenses while maintaining service quality.
	Procurement Path	The ease of procuring the tool by the agency.
Vendor Ecosystem & Strategy	Ecosystem & Support	The ability to build an ecosystem that meets all requirements and have the applicable team to support that ecosystem.
	Geographic Coverage	The data locations that hyperscalers hold their data centers.
	Innovation & Roadmap	The vendor's ability to innovate based on agency needs and plan a roadmap or a portfolio of projects to execute on that innovation.

Scoring methodology

Each strategic theme was assigned a weighted score based on its importance to OSPI's priorities. Vendors were then scored against each theme on a scale from 1-5. Each vendor's weighted scores indicate capability in that strategic theme and its alignment to OSPI's priorities.

- Weighting methodology
  - **0.1 Addition** = Preferred Requirement
  - **0.3 Addition** = Required Requirement
  - **0.5 Addition** = Mandatory Requirement
- Specific Themes Scoring (1-5):
  1. **No Capability.** The capability does not exist in the current architecture, and OSPI cannot build a solution to meet the requirement with current staff.
  2. **Buildable Partial Capability.** The capability does not exist in the current architecture, but OSPI can build a solution that partially solves the requirement (at no additional cost) with current staff.

3. **Buildable Full Capability.** The capability does not exist in the current architecture, but OSPI can build a solution that fully solves the requirement (at no additional cost) with current staff.
4. **Existing Partial Capability.** The capability exists in the current architecture but only partially solves the requirement (may require enhancements or customizations at no additional cost) with current staff.
5. **Existing Full Capability.** The capability exists in the current architecture and fully solves the requirement without the need for modifications.

Microsoft Assessment

Key Strengths

Microsoft maintains specific areas of excellence within its ecosystem, achieving perfect scores in procurement path, data visualization through Power BI integration, compliance/governance frameworks, and database compatibility (SQL Server). The platform also demonstrates solid capabilities in storage flexibility and migration tooling (ExpressRoute), leveraging Azure's mature infrastructure services. These strengths cluster around Microsoft's established enterprise offerings and existing OSPI investments.

Strategic Advantages

1. Procurement Path: Microsoft maintains the strongest acquisition position through existing state contracts and established vendor relationships, enabling the fastest procurement timeline among all options.
2. Data Visualization: The native integration with Power BI provides industry-leading visualization capabilities that are immediately accessible and familiar to existing staff.
3. Compliance & Governance: Microsoft delivers comprehensive compliance frameworks with established Azure governance tools and certifications that align perfectly with current state requirements.
4. Database Compatibility: The platform achieves perfect compatibility with OSPI's existing SQL Server infrastructure, eliminating complex migration challenges and preserving current investments.
5. Storage Flexibility: Azure provides strong storage flexibility through multiple tier options that can accommodate various data types and access patterns.
6. Migration Tooling: Microsoft offers robust migration tools including ExpressRoute and Azure Migration services that facilitate data and application transitions.
7. Geographic Coverage: Microsoft operates from many geographic locations relevant to OSPI's needs (e.g., West US), and can support multi-region deployment, if needed.

Strategic Theme	Weight	Microsoft Full Build	Microsoft Full Build Weighted
Procurement Path	0.10	5	5.10

Strategic Theme	Weight	Microsoft Full Build	Microsoft Full Build Weighted
Data Visualization	0.10	5	5.10
Compliance & Governance	0.10	5	5.10
Database Compatibility	0.10	5	5.10
Storage Flexibility	0.10	4	4.10
Migration Tooling	0.10	4	4.10
Geographic Coverage	0.05	4	4.05

Operational Excellence Areas

1. Interoperability: Microsoft provides moderate interoperability capabilities through Azure integration services, though custom development would be required for comprehensive system-to-system communication.
2. Support & Maintenance: The platform offers standard Azure support channels that improve upon current state but still require significant internal team involvement.
3. High Availability: Microsoft delivers standard high availability configurations that meet basic continuity requirements through Azure's infrastructure.
4. Disaster Recovery: The platform provides adequate disaster recovery capabilities through Azure Site Recovery and backup services.
5. Ecosystem & Support: Microsoft offers access to a moderate partner ecosystem through the Azure marketplace and Microsoft partner network.

Strategic Theme	Weight	Microsoft Full Build	Microsoft Full Build Weighted
Interoperability	0.15	3	3.15
Support & Maintenance	0.10	3	3.10
High Availability	0.10	3	3.10
Disaster Recovery	0.10	3	3.10
Ecosystem & Support	0.10	3	3.10

Considerations

1. Accessibility: Microsoft would require OSPI to build accessibility features from scratch, bearing full operational burden for compliance with accessibility standards.
2. Usability: The platform demands significant technical expertise for operation, resulting in extended time-to-value compared to modern platform alternatives.

3. Manual Work Automations: Only partial automation is achievable through custom development, requiring extensive staff resources and ongoing maintenance.
4. Master Data Management: Basic MDM capabilities would need to be built from scratch with limited current resources and expertise.
5. Managed Services: OSPI would largely self-support custom builds, requiring significant expansion of current team or reliance on external contractors.
6. Compute Fit: The platform requires significant customization for optimal compute configuration, suggesting a 3– 5-year implementation timeline.
7. Cost Optimization: While leveraging existing knowledge, the extensive custom development required would significantly increase total cost of ownership.
8. Innovation & Roadmap: The platform shows limited innovation pathways, requiring custom development for emerging capabilities rather than platform-delivered features.
9. Governance by Design: While compliance tools exist, embedding automated governance into data flows requires extensive custom work beyond current capabilities.
10. Zero-touch Operations: Microsoft has no existing platform capability for automation, requiring complete custom development beyond current staff abilities.
11. Self-Service Enablement: The Microsoft approach provides limited capability for empowering non-technical users without building those capabilities from scratch.

Strategic Theme	Weight	Microsoft Full Build	Microsoft Full Build Weighted
Usability	0.10	2	2.10
Accessibility	0.15	3	3.15
Manual Work Automations	0.15	2	2.15
Master Data Management	0.15	2	2.15
Managed Services	0.10	3	3.10
Compute Fit	0.15	3	3.15
Cost Optimization	0.10	2	2.10
Innovation & Roadmap	0.05	2	2.05
Zero-touch Operations	0.15	1	1.15
Governance by Design	0.10	2	2.10
Self-Service Enablement	0.05	1	1.05

Amazon Web Services + Cloudwick Amorphic Assessment

Key Strengths

The Amorphic platform on AWS demonstrates exceptional capability across critical modernization areas, scoring perfect marks (5) in 15 of the 23 categories. Most notably, it excels in automation capabilities which directly address OSPI's need to reduce operational overhead. The platform's existing full capabilities for interoperability, managed services, and migration tooling suggest a mature, battle-tested solution that could deliver immediate value with minimal customization.

Strategic Advantages

1. Interoperability: Amorphic delivers comprehensive platform interoperability with native connectors, APIs, and data exchange protocols enabling seamless system integration.
2. Usability: The platform provides an intuitive, user-friendly interface designed for both technical and business users, dramatically reducing training requirements.
3. Support & Maintenance: Amorphic offers comprehensive 24/7 managed support with proactive monitoring and maintenance included in the platform.
4. Manual Work Automations: Pre-built automation frameworks immediately eliminate manual processes across data pipelines and operational workflows.
5. Master Data Management: The platform includes enterprise-grade MDM with data lineage, quality management, and golden record capabilities out-of-the-box.
6. Compliance & Governance: Amorphic achieves excellence with comprehensive compliance frameworks including HIPAA, SOC2, and FedRAMP-ready certifications.
7. Compute Fit: The platform automatically optimizes compute resources based on workload patterns, ensuring right-sized infrastructure without manual intervention.
8. Storage Flexibility: Amorphic delivers maximum flexibility with multi-tier, multi-format storage supporting all data types through AWS's extensive storage services.
9. High Availability: The platform achieves maximum uptime through AWS's robust infrastructure with automated failover and redundancy.
10. Disaster Recovery: Amorphic provides comprehensive automated disaster recovery with rapid recovery time objectives and minimal data loss.
11. Ecosystem & Support: Extensive AWS marketplace integration provides access to thousands of pre-integrated tools and services.
12. Geographic Coverage: The platform leverages AWS's global infrastructure perfectly, enabling multi-region deployment and data residency compliance.
13. Innovation & Roadmap: Amorphic continuously delivers platform innovations with integrated AI/ML capabilities and emerging technology adoption.
14. Zero-touch Operations: The platform delivers full zero-touch operations immediately, eliminating manual intervention for routine tasks.

15. Governance by Design: Amorphic embeds governance throughout the platform automatically, ensuring compliance without additional configuration.

Strategic Theme	Weight	AWS+Amorphic	AWS+Amorphic Weighted
Interoperability	0.15	5	5.15
Usability	0.10	5	5.10
Support & Maintenance	0.10	5	5.10
Manual Work Automations	0.15	5	5.15
Master Data Management	0.15	5	5.15
Compliance & Governance	0.10	5	5.10
Compute Fit	0.15	5	5.15
Storage Flexibility	0.10	5	5.10
High Availability	0.10	5	5.10
Disaster Recovery	0.10	5	5.10
Ecosystem & Support	0.10	5	5.10
Geographic Coverage	0.05	5	5.05
Innovation & Roadmap	0.05	5	5.05
Zero-touch Operations	0.15	5	5.15
Governance by Design	0.10	5	5.10

Operational Excellence Areas

1. Accessibility: Amorphic provides strong accessibility features that embed standards directly into the platform, though some customization may enhance compliance.
2. Data Visualization: Amorphic includes robust built-in visualization tools through their frontend applications and built-in medallion capabilities reduce the time needed to get gold data into visualization tools such as Power BI or Tableau.
3. Managed Services: The platform provides substantial managed service capabilities covering most operational needs with some customer involvement.
4. Migration Tooling: The platform provides robust migration toolsets through AWS Database Migration Service and related tools that could bridge the gap between two different providers.
5. Self-Service Enablement: Amorphic enables good self-service capabilities for most user types with intuitive interfaces and guided workflows.

6. Cost Optimization: With most needed capabilities available out of the box, AWS with Cloudwick presents an efficient cost model that also delivers value quickly and efficiently; a slight step behind GCP with Onix in robust additional dedicated support while still better than Microsoft.

Strategic Theme	Weight	AWS+Amorphic	AWS+Amorphic Weighted
Accessibility	0.15	4	4.15
Data Visualization	0.10	4	4.10
Managed Services	0.10	4	4.10
Migration Tooling	0.10	4	4.10
Cost Optimization	0.10	3	3.10
Self-Service Enablement	0.05	4	4.05

Considerations

1. Procurement Path: Procurement will take longer with new complex contracting requirements adding on another vendor to OSPI's list
2. Database Compatibility: Current SQL Server databases may need additional steps to be compatible with AWS storage solution.

Strategic Theme	Weight	AWS+Amorphic	AWS+Amorphic Weighted
Procurement Path	0.10	4	4.10
Database Compatibility	0.10	4	4.10

Google Cloud Provider + Onix Wingspan Assessment

Key Strengths

Wingspan emerges as a high-scoring platform with exceptional performance across critical modernization areas, achieving perfect scores (5) in eleven categories. The platform excels in all three core architectural principles OSPI seeks: zero-touch operations, governance by design, and self-service enablement. Notably, Onix's commitment to providing a dedicated support team and comprehensive training program addresses a critical gap in competing platforms, potentially accelerating adoption and reducing implementation risk.

Strategic Advantages

1. Interoperability: Wingspan provides full platform interoperability with comprehensive APIs and native integration capabilities across diverse systems.
2. Usability: The platform delivers exceptional user experience with intuitive interfaces suitable for all skill levels across the organization.

3. Support & Maintenance: Wingspan offers comprehensive fully managed support with proactive monitoring and complete operational coverage.
4. Manual Work Automations: The platform includes sophisticated automation frameworks that eliminate manual processes across all data operations.
5. Master Data Management: Wingspan provides enterprise-grade MDM solutions with advanced data governance and quality management features.
6. Managed Services: The platform delivers the most comprehensive fully managed experience, minimizing internal operational burden.
7. Compute Fit: Wingspan automatically optimizes compute resources with intelligent workload management ensuring optimal performance.
8. Storage Flexibility: The platform provides maximum storage flexibility through GCP's diverse storage options and intelligent tiering.
9. Cost Optimization: Wingspan achieves exceptional cost optimization through efficient resource utilization and intelligent scaling capabilities.
10. Innovation & Roadmap: The platform continuously integrates cutting-edge technologies including advanced AI/ML capabilities and emerging cloud services.
11. Self-Service Enablement: Wingspan achieves maximum self-service empowerment, enabling all user types to work independently without IT intervention.

Strategic Theme	Weight	GCP+Wingspan	GCP+Wingspan Weighted
Interoperability	0.15	5	5.15
Usability	0.10	5	5.10
Support & Maintenance	0.10	5	5.10
Manual Work Automations	0.15	5	5.15
Master Data Management	0.15	5	5.15
Managed Services	0.10	5	5.10
Compute Fit	0.15	5	5.15
Storage Flexibility	0.10	5	5.10
Cost Optimization	0.10	5	5.10
Innovation & Roadmap	0.05	5	5.05
Self-Service Enablement	0.05	5	5.05

Operational Excellence Areas

1. Accessibility: The platform provides strong accessibility features with built-in compliance tools meeting most regulatory requirements.
2. Data Visualization: Wingspan includes capable visualization tools that deliver strong analytical capabilities for most use cases.
3. Compliance & Governance: The platform offers strong compliance frameworks with comprehensive certifications, though slightly less extensive than competitors.
4. High Availability: Wingspan provides strong high availability architecture through GCP's infrastructure with robust failover mechanisms.
5. Disaster Recovery: The platform delivers strong disaster recovery capabilities with automated backup and recovery procedures.
6. Ecosystem & Support: Wingspan provides access to GCP's growing ecosystem with strong integration options for common enterprise tools.
7. Geographic Coverage: The platform leverages GCP's strong geographic footprint providing good coverage for most deployment scenarios.
8. Zero-touch Operations: Wingspan achieves near-complete automation with minimal manual oversight required for exceptional cases.
9. Governance by Design: The platform embeds strong governance capabilities throughout with some configuration needed for specific requirements.

Strategic Theme	Weight	GCP+Wingspan	GCP+Wingspan Weighted
Accessibility	0.15	4	4.15
Data Visualization	0.10	4	4.10
Compliance & Governance	0.10	4	4.10
High Availability	0.10	4	4.10
Disaster Recovery	0.10	4	4.10
Ecosystem & Support	0.10	4	4.10
Geographic Coverage	0.05	4	4.05
Zero-touch Operations	0.15	4	4.15
Governance by Design	0.10	4	4.10

Considerations

1. Procurement Path: Wingspan faces the longest acquisition timeline through GCP channels over that of Azure and AWS, requiring new vendor relationships and contract negotiations.
2. Database Compatibility: The platform requires more complex database migration strategies, particularly for SQL Server workloads during parallel run.
3. Migration Tooling: GCP's migration tools are less mature for complex enterprise migrations from legacy SQL Server architecture, potentially extending transition timelines.

Strategic Theme	Weight	GCP+Wingspan	GCP+Wingspan Weighted
Procurement Path	0.10	3	3.10
Database Compatibility	0.10	3	3.10
Migration Tooling	0.10	3	3.10

The 36-point scoring gap between the leading platform (Amorphic at 108.5) and Microsoft's native build (72.5) represents a fundamental difference in both approaches outlined in the previous section. Microsoft scored low in critical areas including zero-touch operations and self-service enablement. These gaps cannot be addressed with current OSPI staff capabilities and would require extensive custom development that the platform solutions provide out-of-the-box.

Between the two platforms, AWS Amorphic offers technical capabilities that more closely aligns with OSPI's needs with perfect scores in 15 categories including high availability, disaster recovery, and governance. GCP Wingspan provides strong capabilities with 11 perfect scores while uniquely offering better cost optimization and comprehensive training/support services through Onix's dedicated team. Both platforms transform OSPI's manual, siloed operations into automated, integrated systems that enable data-driven decision-making across all programs.

The financial analysis indicates that while Microsoft has procurement advantages through existing contracts, its cost combined with extensive development requirements makes it the most expensive option long-term. The platforms' higher automation and managed service capabilities reduce operational costs and time-to-value compared to custom build.

Vendor Strategic Assessment Summary Graphic

Strategic Theme	Weight	OSPI Current State Weighted	Microsoft Azure Weighted	AWS + Amorphic Weighted	GCP + Wingspan Weighted
Interoperability	0.15	1.15	3.15	5.15	5.15
Usability	0.10	2.10	2.10	5.10	5.10
Accessibility	0.15	1.15	3.15	4.15	4.15
Procurement Path	0.10	5.10	5.10	4.10	3.10
Data Visualization	0.10	3.10	5.10	4.10	4.10
Support & Maintenance	0.10	1.10	3.10	5.10	5.10
Manual Work Automations	0.15	2.15	2.15	5.15	5.15
Master Data Management	0.15	2.15	2.15	5.15	5.15
Managed Services	0.10	1.10	3.10	4.10	5.10
Compliance & Governance	0.10	2.10	5.10	5.10	4.10
Compute Fit	0.15	2.15	3.15	5.15	5.15
Database Compatibility	0.10	5.10	5.10	4.10	3.10
Storage Flexibility	0.10	1.10	4.10	5.10	5.10
High Availability	0.10	3.10	3.10	5.10	4.10
Disaster Recovery	0.10	2.10	3.10	5.10	4.10
Migration Tooling	0.10	2.10	4.10	4.10	3.10
Cost Optimization	0.10	5.10	2.10	3.10	5.10
Ecosystem & Support	0.10	1.10	3.10	5.10	4.10
Geographic Coverage	0.05	1.05	4.05	5.05	4.05
Innovation & Roadmap	0.05	2.05	2.05	5.05	5.05
Zero-touch Operations	0.15	1.15	1.15	5.15	4.15
Governance by Design	0.10	1.10	2.10	5.10	4.10
Self-Service Enablement	0.05	2.05	1.05	4.05	5.05
	SUMS	**50.5**	**72.5**	**108.5**	**102.5**

**Color scale is relative to other scores within each vendor's column*

7.1.4. Benchmarking Validations

The feasibility study reviewed the experiences of peer agencies who implemented large-scale education data modernization efforts using either a custom-built platform model or a vendor-supported ecosystem. These interviews helped validate assumptions about project timelines, technical complexity, vendor coordination, and staffing requirements. Our findings indicate that OSPI's technical and operational challenges are similar to those experienced by other educational organizations. Many peer agencies have been pursuing modernization for several years and continue to refine their approaches. Key takeaways are summarized below.

Custom-Build Standalone Model with Microsoft Azure

- **Extended Timelines:** A four-year timeline from planning to go-live proved accurate. The agency noted they would have benefited from starting vendor engagement and data modeling a full year earlier.

- **Heavy Vendor Coordination:** Integration with SIS vendors, especially those unable to produce standard-format data, was cited as the most difficult aspect. Weekly vendor meetings were required to maintain progress.
- **Staffing Requirements:** The internal IT department expanded by over 25%, augmented with contractors to support development, knowledge transfer, and SME backfill. They emphasized the need for structured pairing and co-development.
- **Full Customization and Control:** The approach provided complete alignment to state-specific requirements, integration with existing systems, and full ownership of code and architecture without vendor lock-in.
- **High Complexity:** The team built their own DevOps, data pipelines, and Azure infrastructure. They reported that early development was slowed due to lack of production-quality historical data to validate pipelines.

Vendor Platform Model with Amazon Web Services and Cloudwick

- **Rapid Deployment:** The agency implemented their platform within 12 months by leveraging pre-built capabilities and focusing on standards alignment and API integration from day one.
- **All-in-One Ecosystem:** A unified platform included ETL/ELT tools, data cataloging, analytics, and integration gateways. The team also began experimenting with AI/ML using built-in services.
- **Predictable Costs:** Transitioning away from on-prem infrastructure yielded substantial cost savings and enabled clearer budgeting through a SaaS pricing model.
- **Lean Operations:** The shift reduced reliance on traditional DBA roles. IT was reorganized to focus on architecture, operations, and innovation, supported by external partners and in-house R&D staff.
- **Clear Governance Model:** Data stewardship was assigned to business units, not IT. A data governance committee facilitated standards and ownership. Master Data Management and cataloging tools were key enablers.
- **Management Simplification:** The agency mitigated platform complexity by layering partner-built management tools that allowed non-experts to configure integrations and automate processes through a simplified interface.

Benchmarking Summary

These experiences demonstrate the tradeoffs OSPI must weigh in its own modernization effort. Custom-built models offer great flexibility and system ownership but come with higher resource demands and longer timelines. Vendor-supported cloud platforms accelerate delivery and simplify operations but require more alignment to the vendor's architecture and tooling. Either approach will require strong governance from OSPI, early vendor engagement effort, and adequate staffing to succeed.

7.1.5. Approach and Vendor Assessment Summary

This section bridges OSPI's strategic approach analysis with vendor capabilities and readiness, summarizing key lessons learned and narrowing the field of viable delivery models for OSPI's data modernization effort.

A custom-built approach gives OSPI full ownership and flexibility. It also demands significant internal capacity, longer timelines, and sustained vendor coordination. Benchmarking confirms this model is viable, but difficult to execute without strong technical leadership and support.

Vendor-supported platforms offer speed, predictable costs, and a mature toolset. They reduce complexity but require OSPI to operate within the boundaries of the vendor's architecture and roadmap. Success in this model depends on the strength of the partner relationship and the agency's ability to manage adoption across teams.

A hybrid model offers a middle path combining reusable vendor components with OSPI-owned services to maintain flexibility. It allows for faster delivery without giving up control where it matters most.

Each path is feasible. The right one depends on how OSPI balances speed, flexibility, staffing, and procurement. This summary is intended to inform that decision and frame the considerations that will shape the implementation plan that follows this feasibility study.

SECTION 8 – CONFORMITY WITH AGENCY IT PORTFOLIO

This data platform modernization initiative directly implements the Strategic Plan's vision of "Efficient government through technology" by transforming fragmented legacy systems into a unified, cloud-native platform that delivers "Personalized & Optimized Government" services. Built on the strategic values of Partnership, Transparency, and Agility, this transformation operationalizes each of the plan's four strategic pillars through concrete technical implementations and measurable outcomes.

8.1.1. Strategic Focus (Business and IT Goals)

The initiative creates operational efficiency through three key architectural principles—Zero-Touch Operations, Governance by Design, and Self-Service Enablement—that fundamentally transform how OSPI manages and delivers data services. Zero-Touch Operations eliminates the operational burden that typically consumes 70-80% of data team resources, while Governance by Design ensures that security, quality, and compliance are seamlessly woven into every data interaction. Self-Service Enablement democratizes data access across the organization, transforming data from a technical asset managed by IT into a business asset leveraged by everyone.

The proposed solution implements a seven-layer architecture stack built on medallion architecture (Bronze/Silver/Gold zones) that provides progressive data refinement with clear quality levels. This approach addresses OSPI's most prevalent operational issues identified in the Pareto Analysis, including missing or improperly structured operational processes, manual data processing, insufficient data quality standards, and limited analytics access. Through an agile modernization strategy built on rapid learnings and end-to-end releases, each phase delivers standalone value while building toward comprehensive transformation.

8.1.2. Effect on Technology Infrastructure

Pillar 1: Efficient & Effective Government

Creating Opportunities for Operational Efficiency & Improving Constituent Access to Services

The platform eliminates redundant infrastructure by implementing the proposed seven-layer architecture stack that consolidates disparate data systems into a unified cloud-native solution. The Governance Control Plane serves as a central hub with policy automation engines enforcing rules across all data assets, continuous compliance validation with automated remediation, and centralized audit aggregation. This consolidation directly addresses OSPI's fragmented SQL Server systems and decades of technical debt through common, reusable components that reduce infrastructure complexity while improving service delivery efficiency.

Create Constituent-Focused Interfaces

Self-Service Enablement principles provide frictionless data discovery where business users can browse, search, and request access to datasets through intuitive data catalogs without submitting IT tickets. Natural language processing powers conversational analytics enabling users to ask questions in plain English and receive immediate visualizations. The Serving Layer, on a simpler level, would be able to provide governed and approved data directly accessible through current methods of data dashboards and file downloads, essentially even without changing interfaces drastically, end-users would already be experiencing benefits through better data with their current interfaces.

Develop Accessible & Responsive Designs

The platform's user-friendly interfaces built from Gold-layer data enable program managers to build reporting solutions independently without requiring programming expertise. Every interface implements responsive design patterns with automated report builders that reduce time-to-insight from weeks to hours. The semantic business layer provides pre-built metrics, calculated fields, and standardized business definitions ensuring consistent interpretation across all consumer touchpoints and reporting interfaces.

Increase Access to Open Data

The Governance Components layer implements automated data cataloging with business and technical metadata, making all non-sensitive data discoverable. Visual lineage tracking from source to consumption provides complete transparency, while dynamic access controls with attribute-based policies maximize data availability. The platform's automated metadata capture instantly captures every dataset, creating a searchable knowledge base that serves as the single source of truth for organizational data assets.

Pillar 2: IT Workforce

Recruit, develop & retain exceptional technology staff & leaders

Champion Competitive Job Classification Structure

Zero-Touch Operations eliminates the operational burden consuming 70-80% of data team resources, freeing staff to focus on innovation rather than maintenance. Event-driven orchestration enables pipelines to activate instantly upon data arrival, while policy inheritance and propagation mechanisms cascade automatically through data lineage. This automation makes positions more attractive by eliminating mundane maintenance tasks and enabling focus on strategic initiatives.

Develop Enterprise Workforce Development Strategy

The phased migration approach starts with small, high-value data projects, treating each step as a standalone release delivering end-to-end value. Capability clarifies skill requirements per role, allowing OSPI to design targeted training aligned to platform usage and governance responsibilities. The platform includes development of reusable components for a Feature Store—a centralized library of cleaned and engineered data features—that can be templated and used repeatedly.

Re-imagine Management Practices for Employee Empowerment

Infrastructure resources benefit from dynamic optimization with intelligent orchestration including dependency management and automatic retries. Continuous quality monitoring operates through real-time data profiling with anomaly detection, completeness checks, and business rule validation triggering proactive alerts. The platform provides FinOps integration for cost attribution and optimization, empowering teams with transparency into resource usage and enabling data-driven operational decisions.

Pillar 3: Enterprise Architecture

Identify strategic technology investments to support common business needs/functions.

Leverage Cloud Solutions

The cloud-agnostic modern data platform architecture provides a proven blueprint for building enterprise-scale data solutions across any major cloud provider (Azure, AWS, or Google Cloud). This "Automation First" design philosophy drives every architectural decision, creating self-operating, self-healing, and self-governing systems that eliminate traditional bottlenecks where manual data management creates delays and requires extensive teams.

Modernize Infrastructure and Applications

The Processing Layer implements elastic compute engines for distributed batch and stream processing with serverless functions triggering on events without infrastructure management. The platform scales elastically based on real-time workload patterns, automatically provisioning compute capacity during peak demand and deallocating resources during idle times. Migration thresholds activate at large data volumes or when batch jobs exceed four hours, ensuring optimal performance at any scale.

Increase Capacity to Manage & Share Information

The medallion architecture's Storage Layer provides progressive data refinement: Bronze preserves immutable raw data, Silver contains cleansed and validated datasets, and Gold delivers business-ready aggregations and feature stores. End-to-end lineage tracking provides complete data journey visualization from source systems through all transformations to final reports, enabling comprehensive impact analysis and rapid root cause investigation.

Evaluate Options for Shared Solutions

The platform's multi-pattern data capture in the Ingestion Layer supports batch, streaming, CDC, APIs, and file drops with inline validation and automatic quality checks. Standardized data contracts and API gateways facilitate ecosystem integration, while the architecture supports potential expansion to serve other state agencies or school districts through isolated namespaces with resource tagging and chargeback capabilities for cost allocation.

Pillar 4: Security & Privacy

Find and stop security risks while increasing privacy awareness

Proactive Security Architecture

The Foundational Infrastructure layer establishes a secure cloud landing zone with isolated VPCs/vNets and private endpoints, centralized identity management through SSO and role-based access control, with potential for granular controls down to the data table row-level or column-level access. Infrastructure as Code enforces security policies consistently across all environments with unified observability through centralized logging, monitoring, and compliance scanning. This directly addresses OSPI's compliance and security challenges identified as the ninth most prevalent issue category in the Pareto Analysis.

Continuously Improve State Defenses

Zero-trust security model implementation provides policy-driven access controls with encryption at rest and in transit, comprehensive audit logging, and automated compliance reporting built directly into the data flow. The platform implements defense-in-depth with enterprise encryption using managed key vaults and automatic rotation, while the Governance Control Plane maintains continuous compliance validation with automated remediation and intelligent alerting with context-aware incident management.

Develop Systems with Protection and Privacy Focus

Governance by Design principle embeds Privacy by Design through automated sensitive data discovery and classification at the Ingestion Layer. The platform automatically identifies and tags sensitive information including personally identifiable information and protected health information. Smart error handling uses quarantine zones and retry logic while sensitive data isolation provides automatic PII/PHI detection and masking, ensuring appropriate handling and compliance with FERPA, COPPA, and state privacy regulations.

Provide Staff with Tools and Training

The platform includes isolated training environments for safe practice with security tools and incident response procedures. Real-time governance dashboards enable proactive compliance monitoring, flagging anomalies that contradict compliance rules and instantly alerting data stewards and owners. The infrastructure provides comprehensive audit capabilities through centralized audit aggregation for forensics and regulatory reporting.

SECTION 9 – PROJECT MANAGEMENT AND ORGANIZATION

The success of OSPI's data modernization initiative will rely on a clearly defined and accountable project management structure, sustained leadership engagement, and disciplined execution. This section outlines the proposed approach for managing the implementation effort covering team structure, roles and responsibilities, decision-making, governance, quality assurance, and organizational alignment.

9.1.1. Project Charter and Governance Structure

This initiative will be governed by a formal Project Charter that defines the project's purpose, scope, guiding principles, and delivery structure. The charter will:

- Establish delivery accountability across OSPI and vendor partners
- Define escalation paths, decision-making boundaries, and risk management processes
- Be approved by OSPI executive leadership and the Steering Committee
- Apply to both the implementation phase and any bridge work that precedes it

A three-tiered governance framework will provide oversight and decision authority:

- **Executive Steering Committee (ESC):** Sets strategic direction, manages project scope, schedule, and budget, and resolves escalated issues. The Executive Sponsor retains ultimate authority.
- **Project Manager:** Central point of contact for day-to-day decisions, responsible for status reporting, escalation, and issue resolution.
- **Project Liaison Team:** A working team that includes representation from all stakeholders (e.g., Cybersecurity, Enterprise Architecture), empowered to resolve issues quickly and escalate when needed.

This structure reinforces accountability, enables rapid delivery, and ensures alignment with OSPI's strategic plan.

9.1.2. Project Schedule and Work Plan

The modernization effort will follow a multi-year phased implementation, aligned to WaTech's seven-year cost modeling requirements. The first 12–18 months will focus on platform stand-up, governance framework development, and product team formation. Migration activities will be sequenced by domain, prioritized by value and complexity.

A detailed migration timeline is outlined in Section 10. Core modernization activities are expected to conclude within five years, with optimization and sustainment extending through year seven.

9.1.3. Program Management Office (PMO)

A centralized Program Management Office will coordinate across all workstreams, manage delivery velocity, and serve as the single point of coordination between OSPI, WaTech, vendors, and internal stakeholders. The PMO will operate under a hybrid methodology:

- Agile delivery at the team level, including sprint reviews and iterative delivery
- Structured program-level oversight for tracking progress, risks, dependencies, and alignment

Key PMO functions will include:

- RAID (Risk, Action, Issue, Decision) log management
- Budget/schedule/scope change tracking
- Vendor integration and performance monitoring
- Progress reporting aligned to Executive Milestones and legislative cycles
- Change management coordination

The PMO will also ensure compliance with Section 701 of the Washington State budget, which requires:

- Agile delivery with live system demos every two weeks
- Initial product deployment within 180 days of contract execution
- Retention of critical project functions by state personnel

9.1.4. Decision-Making Process

OSPI's decision-making framework is grounded in its Enterprise Architecture Charter and designed to balance flexibility with consistency. The model supports quick execution while ensuring alignment with long-term goals.

Key Governance Bodies

- **Architecture Review Board (ARB):** Meets monthly to review design proposals, promote reuse of patterns, and ensure consistency with enterprise architecture principles. Reviews are collaborative, not evaluative.
- **Change Advisory Board (CAB):** Oversees system changes, combining automation for routine changes with human oversight for complex ones. Emergency CAB processes allow convening within 30 minutes when needed.
- **FinOps Council:** Joint finance-engineering group that monitors cloud spend, enforces tagging policies, and aligns investments with budget/funding sources.
- **Data Governance Council:** Oversees FERPA compliance, data classification, and master data management standards. Ensures all data products align with medallion architecture standards (Bronze/Silver/Gold).

Decision Criteria

All major decisions will be reviewed against five core architectural principles:

1. Start small, improve iteratively
2. Design for interoperability and reuse
3. Prioritize simplicity and sustainability
4. Build on open standards and secure foundations
5. Support autonomy with guardrails

Architecture Decision Records (ADRs):

All major technology decisions will be documented using a standardized ADR format that includes:

- Context and problem statement
- Alternatives considered
- Chosen approach and rationale
- Implementation implications
- Review and signoff from relevant stakeholders

ADRs will be stored in a centralized repository (SharePoint or GitHub) for future reference and reuse.

9.1.5. Product Team Organization

To better support agile delivery and continuous analytics development, OSPI would adopt a product-oriented organizational model. This model replaces siloed functional teams with eight permanent data product teams, each with end-to-end responsibility for its domain.

Each product team will:

- Own its release cadence and backlog
- Operate independently but coordinate through a shared platform and governance structure
- Deliver complete product lifecycle management: development, testing, production, maintenance, continuous improvement

The Compliance Data Product Team will serve as a shared governance backbone—enforcing standards, access controls, naming conventions, and metadata frameworks.

Product teams will be supported by cross-cutting functions:

- Platform & Cloud Engineers – maintain underlying infrastructure and automation
- Solution Architects – align implementation to enterprise architecture standards
- Change & Training Leads – support adoption and workforce readiness

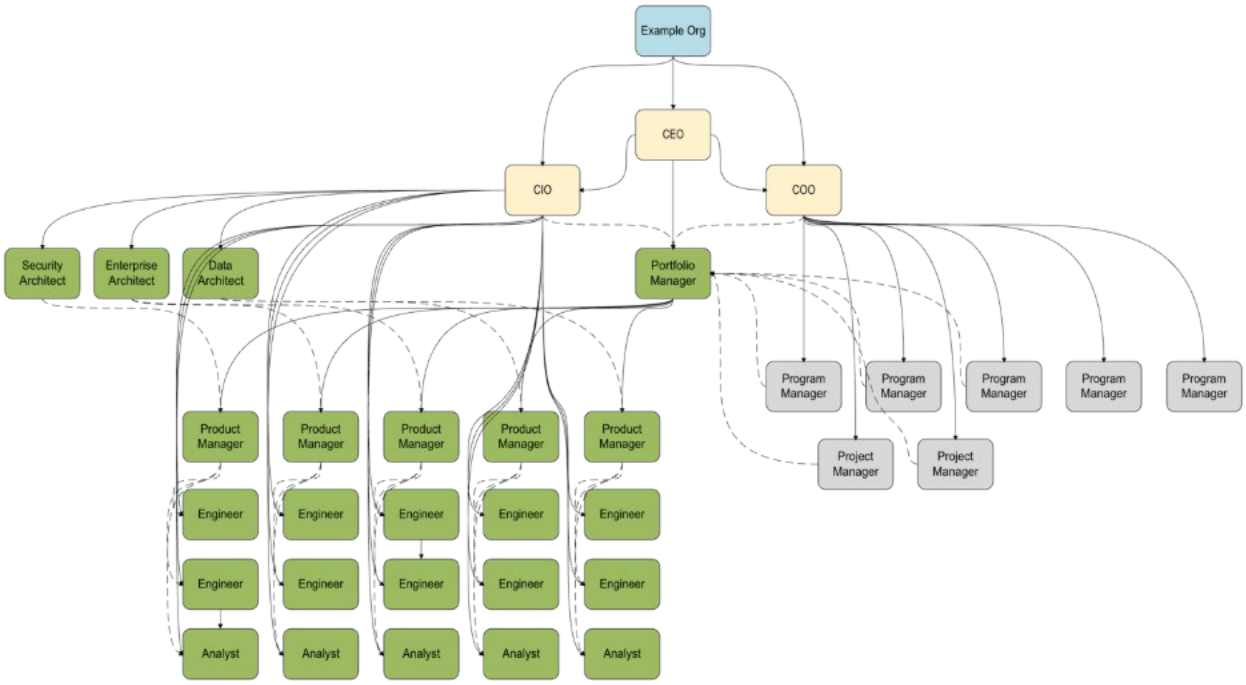

Example Product-Centric Organization Diagram

9.1.6. Key Leadership Qualifications

Leadership across this program will require fluency in both strategy and execution, alongside delivery experience in the public sector. Key qualifications for executives, product owners, project managers, and technical leads include:

- **Strategic Alignment:** Ability to connect delivery with OSPI's mission, equity priorities, and policy drivers
- **Cross-Functional Coordination:** Managing diverse teams that span business and technical roles
- **Governance & Compliance Literacy:** Familiarity with privacy, metadata, and security frameworks
- **Vendor Management:** Coordinating partners, enforcing accountability, and managing service-level performance
- **Public Sector Navigation**: Understanding of Washington's budget processes, procurement structures, and regulatory requirements
- **Change Navigation:** Guiding teams through role transitions, tooling shifts, and adoption of modern processes

Leadership responsibilities will be distributed across OSPI executives, program sponsors, technical managers, and vendor partners based on the phase, data domain, and system scope for each project iteration.

9.1.7. Quality Assurance Strategies

Due to the complexity and risk associated with this project, external Quality Assurance (QA) will be required per WaTech standards. The QA provider will conduct independent assessments of:

- Project readiness
- Implementation quality
- Risk mitigation plans
- Organizational adoption progress

In addition:

- Clear deliverables and expectations will be defined during contracting
- Frequent system demos will validate functionality and reduce rework
- OSPI staff will be actively involved in testing, issue identification, and resolution
- Feedback loops will ensure expectations are being met early and often

The PMO and Project Manager will jointly maintain real-time visibility for the Steering Committee and ensure transparency throughout the delivery lifecycle.

SECTION 10 – ESTIMATED TIMEFRAME AND WORK PLAN

Since 2023, publicly funded projects implemented in Washington departments are required to deliver on an Agile model. Section 701 outlines the special appropriations language for the current operations budget (ESSB 5187), and requires:

- Agile development methodology with a live system demonstration every two weeks
- Product deployment is due within 180 days of a signed contract
- Key project functions deemed critical must be retained by state personnel and cannot be outsourced

Given the Section 701 mandate, project stakeholders are expected to align and deliver from rapid development cycles. The 180-day deployment timeline will require early and continuous engagement from functional users, dedicated testers, and technical teams to validate functionality in real-time. ISG recommends that OSPI negotiate delivery terms during procurement efforts that align to Section 701 requirements while also allowing for a sustainable, phased rollout across districts.

10.1.1. Data Migration

Standalone Release Strategy

OSPI's data migration adopts an incremental release approach where each release delivers production-ready, end-to-end functionality that provides immediate value. Rather than waiting for a complete platform transformation, this strategy ensures that individual processes and teams can leverage cloud capabilities throughout the 5–7-year journey to becoming fully cloud native. Each release represents a complete source-to-output data flow that can operate independently while building toward the comprehensive platform vision.

Migration Readiness & Prerequisites

Before initiating any migration release, OSPI must ensure comprehensive preparation:

- **Organizational Readiness Assessment**: Validate that impacted teams have necessary training, resources, and change management support
- **Technical Dependency Mapping**: Document all upstream and downstream systems, APIs, and data flows for each migration candidate
- **Governance Framework Establishment**: Define data ownership, quality standards, and compliance requirements before moving data
- **Infrastructure Foundation**: Deploy core platform components including networking, security, and base storage layers

- **Communication Protocols**: Establish clear channels between data product teams, stakeholders, and consumers

Release 0: Foundation Phase (2-4+ Weeks Per Release) Every data product migration begins with a critical discovery and planning release:

- **Data Discovery & Cataloging**: Use Microsoft Purview to identify all data sources, schemas, and relationships
- **Conflict Resolution**: Standardize competing metric definitions and eliminate redundant databases or tables
- **Stakeholder Validation**: Confirm migration scope and ensure only relevant, active data is included
- **Governance Design**: Define policies, ownership models, quality rules, and compliance requirements
- **Transformation Mapping**: Design the complete Raw → Bronze → Silver → Gold data flow architecture
- **Delivery Preparation**: Determine needs for APIs, dashboards, reporting interfaces, and user access patterns

Step 1	Step 2	Step 3	Step 4	Step 5
On-Prem	ExpressRoute	Azure Data Factory	Purview + Unity Catalog	Staging/Bronze/Silver/Gold
Connect on-prem data to Integration Runtime VM	Configure ExpressRoute Connection	Set up cadence for pipeline runs	Run auto-discovery and automate security	Validate with HITL. Automate transformations in established framework

Migration Data Pipeline Implementation Framework

COMPONENT	PURPOSE	MIGRATION PHASE	POST-MIGRATION
Databricks / Spark-Based Fabric	• Data processing • Orchestration • Medallion architecture • Governance Automation	Required	Required
Purview and/or Unity Catalog	• Data Discovery Automation • Governance Control Plane	Required	Required
ADLS Gen2	• Data lake storage	Required	Required

COMPONENT	PURPOSE	MIGRATION PHASE	POST-MIGRATION
	• Bronze/Silver/Gold layers		
ExpressRoute	• Secure, reliable connectivity • Handle multiple SQL servers	Required	Optional
Integration Runtime VMs	• Extract data from on-prem SQL • 2-4 VMs for load distribution	Required	Decommission
Key Vault	• Secrets & connection strings • Security keys	Required	Required
Virtual Network	• Network isolation • Security boundaries	Required	Required
Azure Monitor	• Logging & alerting • Performance monitoring	Required	Required
Private Link	• Secure private endpoints • No public exposure	Required	Required

Migration Tech Stack (Core)

10.1.2. Migration Template

1. EXTRACT: Set up data ingestion

- CDC/batch extraction from source SQL
- Land in ephemeral staging layer (raw format)

2. HITL: Migrate logic

- Convert SQL stored procedures to Spark/SQL
- Graduate to Bronze layer
- Add data quality checks

3. TRANSFORM: Create Silver layer

- Cleanse datasets, dedupe, fix anomalies
- Organize data into established framework

4. DEMOCRATIZE: Create Gold layer

- Business-ready datasets
- Connect to consumers

5. UNIT TEST: Parallel run period

- Compare outputs

- Performance testing

- User acceptance

6. CUTOVER: Go live

- Redirect applications

- Decommission old process and system

ISG recommends starting with a narrow, high-value system to show results quickly, while incrementally adding the processes and tooling that support sustainable, repeatable cloud modernization across additional data domains.

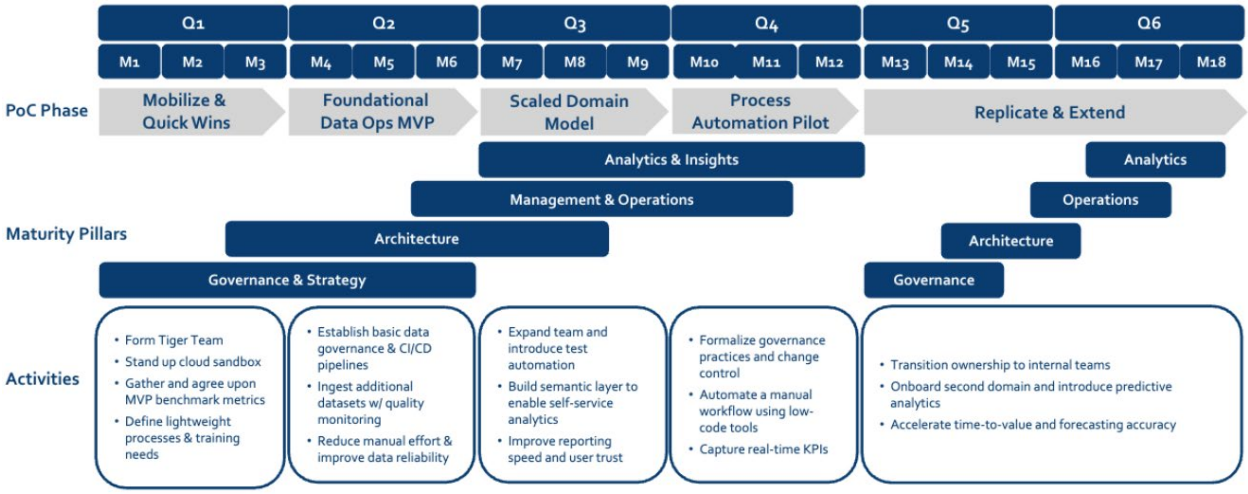

Example Migration Journey: Report Card System

Demonstrating how releases build incrementally while delivering standalone value:

- **Release 1 - Enrollment Data**: Establishes student demographic foundation with full governance, enabling immediate enrollment reporting
- **Release 2 - Assessment Scores**: Adds test data while reusing enrollment infrastructure, delivering assessment analytics
- **Release 3 - Financial Data**: Incorporates budget information leveraging existing student data for per-pupil calculations
- **Release 4 - Educator Data**: Adds staffing information using enrollment for student-teacher ratios
- **Release 5 - Graduation Outcomes**: Integrates all prior data for comprehensive outcome tracking

Result: Each release provides immediate value (enrollment reports, test analytics, etc.) while building toward a fully integrated Report Card system.

Inter-Team Dependencies & Coordination

The Compliance team serves as the automation hub, providing:

- **Governance Templates**: Reusable compliance automation patterns for all data products
- **Security Standards**: Centralized security policies applied across all migrations
- **Quality Frameworks**: Standardized data quality rules and validation logic
- **Legislative Compliance**: Ensuring all migrations meet FERPA, COPPA, and state requirements

Other data product teams integrate these standards while maintaining independence for their specific domain migrations.

10.1.3. Risk Mitigation Strategies

- **Parallel Operations**: Maintain legacy systems during transition with automated comparison reports
- **Incremental Cutovers**: Migrate user groups gradually rather than big-bang transitions
- **Rollback Procedures**: Each release includes documented rollback plans with recovery point objectives
- **Pilot Programs**: Start with non-critical systems or willing early adopters for each data product
- **Vendor Partnerships**: Engage experts and Platform-as-a-Service providers for complex migrations

10.1.4. Success Metrics & Validation

Each release must demonstrate:

- **Functional Completeness**: All required features operational in production
- **Performance Standards**: Meeting or exceeding legacy system response times
- **Quality Thresholds**: X% data accuracy compared to source systems
- **User Adoption**: Defined usage metrics and satisfaction scores
- **Cost Efficiency**: Cloud costs within budgeted parameters
- **Compliance Verification**: Passing all security and regulatory requirements

10.1.5. Estimated Migration Timeline

- **Year 1-3:**
    - Foundation infrastructure and Compliance team standards (enabling all subsequent migrations)
    - Core data products (Student, Finance, Educator) with highest daily usage

- **Year 3-4**: Complex integrations (Assessment, Services) requiring extensive transformation
- **Year 4-5**: Operational systems and remaining legacy applications
- **Year 5+**: Optimization, advanced analytics, and continuous modernization

This approach ensures that while the complete transformation spans 5-7 years, every quarter delivers tangible improvements that teams can immediately leverage, maintaining momentum and demonstrating continuous value throughout the migration journey.

To be fully cloud-native with 100+ TBs of data in siloed processes within 5-7 years will be an aggressive endeavor. The golden triangle of project management will apply if the migration scope, time, or cost is to be altered.

Project Management Golden Triangle

Understanding the Golden Triangle

The Project Management Golden Triangle represents the fundamental relationship between three critical constraints that define the migration's success. These constraints are interconnected – changing one inevitably affects the others.

The Core Principle

You can optimize for two constraints, but the third will be compromised. For example, the migration can quickly deliver high quality, but it will cost more. Or the product outputs can have high quality at low cost, but it will take longer. To optimize for both cost and timeline, OSPI should expect a narrower scope or take advantage of vendor accelerators.

Scope/Quality Time Cost

The features, functionality, and quality standards of the deliverable. What needs to be accomplished and how well it should be done.	The schedule and deadlines for project completion. How quickly the project needs to be delivered.	The budget and resources available for the project. The financial and human resources allocated.

Carefully going through OSPI's desired outcomes using the Decision Framework from Section 7 will help determine the right scope, timing, and cost that makes the most sense for OSPI.

SECTION 11 – COST ANALYSIS

A summary of the estimated costs of this modernization effort is provided below. The summary addresses non-recurring costs for the Implementation phase, as well as recurring annual cost estimates for Maintenance & Operations. The scope of this Cost Benefit Analysis is 7-years per WaTech guidelines, with all major data migration operations concluding within 5 years of project inception. Staffing, infrastructure, and data storage costs are driven by key pillars that emerged from the feasibility study rooted in both current state discovery and the strategic need to transition OSPI from a functionally aligned organization to a product-oriented enterprise.

11.1.1. Key Pillars Driving the Cost-Model

- **Zero-Touch Operations:** Cost projections include investment in automation tools and cloud-native services that reduce manual system intervention and increase reliability. The target state automates pipeline orchestration, monitoring, and scaling through policy-driven tools, reducing manual intervention and enabling reliable operations.
- **Governance by Design:** Rather than treating data governance as a bolt-on compliance function, the future-state architecture embeds governance policies directly into systems, platforms, and work processes. This includes enforcing access controls, lineage tracking, and data quality rules within the platform itself. This ensures compliance and transparency at scale without introducing additional operational drag to OSPI staff.
- **Self-Service Enablement:** The cost model anticipates the creation of role-based access layers and user-friendly tools that allow analysts, program staff, and external partners to explore and use data without relying on central IT functions. Enabling secure and governed access to data products will reduce bottlenecks and empower users to answer questions and inform decisions independently.

Modern data platforms paired with strong governance can deliver returns within the first year, particularly through reduced data preparation time and fewer meetings spent troubleshooting data quality issues. Meaningful return on investment typically materializes within 2-3 years, with significant organizational maturity emerging in year four and beyond for large-scale legacy transformations. These timelines vary considerably based on the organization's readiness to adopt new practices and the availability of team members prepared to take on expanded responsibilities. While some agencies achieve steady-state operations within five years, others - particularly those with extensive technical debt or complex regulatory requirements - may require significantly longer.

The return on investment typically stems from four key areas: operational savings from automation and efficiency gains, risk mitigation through improved compliance and data quality, faster time-to-insight enabling quicker decision-making, and ultimately better business outcomes from data-driven strategies. The reality is that data transformation is a journey requiring sustained commitment rather than a quick win - which is why the release-by-release plan outlined in Section 6 focuses on delivering incremental value throughout the implementation rather than waiting for a single "big bang" payoff.

11.1.2. Key Assumptions Driving the Cost-Model

The following key assumptions drive the feasibility study's cost modeling, especially regarding staffing, infrastructure, and migration velocity.

- The first year of effort is centered around a Proof of Concept (PoC) and migration foundations.
    - The vendor partners with OSPI employees, contractors, and professional services to perform PoC and foundational migration work.
    - OSPI will prove the repeatable migration model through a pilot business segment.
    - OSPI will take the time required to ensure things are fully prepared and communicated for a successful migration (e.g. Cloud environments ready).
- The significant portion of required resources for data modernization and cloud management do not exist at OSPI today and will require hiring.
    - From the feasibility study's discovery data, OSPI is assessed to have 11 of the 28 roles necessary to perform this work to successful completion. The cost model augments all necessary staffing gaps.
    - Gaps are heavily outsourced with Full-Time Contractors (FTCs) to start, then transition from FTCs to Full-Time Employees (FTEs) as OSPI gains traction during implementation.
    - OSPI is short-staffed, even with the roles that are already present. The cost model estimates the headcount gaps required in each role.
- Training and upskilling costs will be required to build long-term internal capability.
    - Budget has been allocated for training current staff, new hires, and leadership on future state architecture.
    - Training costs are estimated to decrease as OSPI matures internal team development operations.
- Cloud platform licensing and service costs will scale with migration progress.
    - The migration assumption is to move roughly 20TBs each year across several projects
    - As data is rebuilt in the cloud, costs will grow from the newly migrated data, and the data previously migrated & maintained in storage.
    - Data volume is assumed to continue growing by 5TB each year with "business-as-usual" operations.
    - In the first several years, operations will run parallel between on-premises processes and cloud processes (dual costs), with eventual sequential cutover and decommissioning.
- Tooling costs (e.g. governance, metadata, CI/CD, ETL) are assumed to follow an "MVP" approach.

- Each year, based on work allocated, only the most necessary tools would be "licensed" or added to the tech stack.
- Some tools will become unnecessary and terminate at the end of the migration (e.g. Integration Runtime VMs)

- The sustainment phase will use a lean operations model with targeted automation.
  - A clear drop of FTCs begins in Year 4 as FTEs are assumed to be trained and ready to take over the data & cloud work.
  - Training should continue through a Learning & Development program to keep skills sharp as OSPI's needs evolve.

11.1.3. Staffing Model

As mentioned in the key assumptions above, the technical staff required for OSPI to conduct a full-scale data migration to the cloud within a 5-year horizon requires a significant investment. This staffing model transforms OSPI from a functionally based organization to a full-stack, product-oriented organization that can scale and deliver enterprise-grade reporting and eventually sustain data-related demands relatively independently of third-party supports. There are existing roles at OSPI that are aligned with various staffing positions in this model, including a Database Administrator and Data Analysts. The feasibility study Cost Analysis includes recommended employee, contractor, and professional services headcount from year 1-7 of the project. This headcount recommendation represents the full estimation of effort required for each role. A skills and capacity assessment is recommended on current OSPI roles relevant to this project to determine the necessary development & support required for existing roles, as well as additional hiring count required for each role against the study's headcount estimations.

The following roles are seen as critical for OSPI's success in data modernization and cloud operations.

Category	Position	Responsibilities	Sourcing
Product & Program Leadership	Executive Sponsor	Champions the vision, secures funding, and resolves executive-level blockers	OSPI FTE
	Business Sponsor	Aligns the business goals, drives stakeholder support, and ensures ROI	OSPI FTE
	Technology Sponsor	Guides technical direction, ensures architecture and platform fit.	OSPI FTE
	Product Owner	Prioritizes backlog and ensures product meets user needs.	OSPI FTE
	Business Analyst	Gathers, documents, and validates business requirements and workflows.	FTC to start, eventual FTE
	Product Manager	Defines product strategy and aligns delivery with business needs.	OSPI FTE

Category	Position	Responsibilities	Sourcing
	Portfolio Manager	Oversees strategic alignment and value delivery across initiatives.	FTC to start, eventual FTE
	Project Manager	Plans, tracks, and manages delivery timelines, scope, and resources.	OSPI FTE
	Program Manager	Coordinates delivery across projects to meet program goals.	FTC to start, eventual FTE

Process & Change Management	**Process Mapping Expert**	**Documents and improved business processes and workflows.**	**FTC terminates at project completion**
	Change Manager	Develops and executes organizational change strategies.	FTC terminates at project completion
	Change Management Analyst / Trainer	Delivers communications and end-user training for adoption.	FTC terminates at project completion
	Independent QA Specialist	Validates solution quality, compliance, and readiness for go-live.	FTC terminates at project completion

Architecture & Design	**Application Architect**	**Designs application components, patterns, and system integration.**	**FTC to start, eventual FTE**
	Data Architect	Designs data models, flows, and governance for enterprise data.	FTC + FTE to start, eventual FTE
	Enterprise Architect	Aligns business, data, and technology strategy in solution design.	OSPI FTE

Engineering & Development	**Junior .NET Developer**	**Builds and maintains application features under senior guidance.**	**FTC terminates at project completion**
	Senior .NET Developer	Designs and delivers scalable .NET solutions aligned with architecture.	FTC terminates at project completion
	QA/Test Engineer	Designs and executes test cases to validate functionality and performance.	FTC terminates at project completion
	Test Automation Engineer	Develops automated test scripts for continuous testing coverage.	FTC + FTE to start, eventual FTE
	Cloud Support Engineer	Manages cloud operations, deployments, and issue resolution.	FTC terminates at project completion
	Junior Cloud Support Engineer	Assists in cloud system monitoring and basic troubleshooting.	FTC terminates at project completion

Security & Infrastructure	**Cybersecurity Architect**	**Designs enterprise security models, controls, and architectures.**	**OSPI FTE**
	Cybersecurity Engineer	Implements and maintains security tools and cyber threat defenses.	FTC + FTE to start, eventual FTE
	Cloud Platform Engineer	Builds and configures scalable, secure cloud environments.	FTC, terminates at project completion
	Database Administrator	Manages database performance, backups, and access controls	OSPI FTE

Data Analytics	**Data Analyst**	**Delivers reports and insights using structured analysis and visualization.**	**OSPI FTE**
	Data Platform Engineer	Builds and manages data platforms and integration pipelines.	FTC, terminates at project completion
	Data Engineer	Designs and develops scalable data pipelines and transformation logic.	FTC + FTE to start, eventual FTE
	BI Engineer	Develops dashboards and reporting solutions for business intelligence.	FTC, terminates at project completion
	Data Scientist	Builds predictive models and advanced analytics to support decisions.	FTC, terminates at project completion
	Data Product Manager	Defines and manages the roadmap for data products and assets.	FTC + FTE to start, eventual FTE

11.1.4. Non-Recurring Implementation Costs

The non-recurring costs represent initial investment required to implement the data modernization effort across infrastructure, development, staffing, and supports. These costs vary between the vendors due to different levels of automation, platform maturity, and support team offerings. Microsoft's model assumes a heavier investment in services, while AWS and GCP leverage more automation and managed capabilities which reduce staffing requirements and implementation overhead.

7-Year Non-Recurring Costs – OSPI Data Modernization			
Cost Category	Microsoft	Amazon Web Services	Google Cloud Provider
Infrastructure, Development	$292.1K	$639.6K	$1.2M
FTC and Professional Services	$87.1M	$47.9M	$48.7M
Total Implementation Costs	**$87.4M**	**$48.5M**	**$49.9M**

11.1.5. Recurring Maintenance & Operations Costs

Recurring M&O costs reflect the ongoing expenditures needed to sustain the modernized environment over the 7-year study window. This includes staffing, software licenses, infrastructure upkeep, and support services. While FTE costs remain constant across vendors, the need for external support and infrastructure spend varies based on each vendor's managed service model.

7-Year Recurring Maintenance & Operations Costs – OSPI Data Modernization			
Cost Category	Microsoft	Amazon Web Services	Google Cloud Provider
Infrastructure, Development	$7.5M	$6M	$2.1M
Software	$0	$3.2M	$2.2M
Service Contracts & Support	$84K	$0	$1.9M
FTE Salary & Wages	$36M	$36M	$36M
Total M&O Costs	**$43.6M**	**$45.2M**	**$42.4M**

11.1.6. Summary Costs

Data modernization implementation costs vary significantly based on organizational size, existing technical debt, and regulatory complexity, ranging from hundreds of thousands for focused initiatives to over $100M for large-scale agency transformations attempting to fully overhaul legacy systems and processes. These investments typically span multiple years and encompass not just technology, but the comprehensive restructuring of how an organization creates, manages, and derives value from its data assets. The total 7-year cost estimate combines both implementation and M&O to provide a thorough view of the long-term investment. Microsoft presents the highest total cost due to its development-heavy delivery model, while AWS and GCP offer lower total costs driven by more platform-based efficiencies and automation through competitive partnerships.

7-Year Costs Summary (Implementation + Maintenance & Operations)			
Cost Category	Microsoft	Amazon Web Services	Google Cloud Provider
Implementation	$87.4M	$48.5M	$49.9M
Maintenance & Operations	$43.6M	$45.2M	$42.4M
7 Year Totals	**$131M**	**$93.7M**	**$92.2M**

SECTION 12 – INCREMENTAL COSTS

No additional incremental costs were identified beyond those already included in the main cost model. OSPI's modernization cost structure assumes full coverage of implementation, staffing, licensing, infrastructure, transition, and training within the established cost categories.

If any additional short-term systems, contracts, or staffing are identified during implementation planning, they can be incorporated into subsequent funding requests or procurement activities. Any temporary tools or stopgap solutions used during implementation should be tracked as part of formal budget planning.

SECTION 13 – BENEFITS

Modernizing OSPI's data environment will deliver transformative benefits across the agency, impacting decision-making, operational efficiency, and stakeholder experience. These benefits extend well beyond technology upgrades positioning OSPI to become a data-driven organization that supports equity and continuous improvement in Washington's education system. Key anticipated benefits include:

Improved Decision-Making Through Data Accessibility

A modern data platform provides real-time, accurate, and integrated data to policymakers, educators, and district leaders. This empowers OSPI to:

- Shift from reactive reporting to proactive, evidence-based decisions.
- Identify trends, equity gaps, and performance indicators faster.
- Enable timely policy interventions and targeted resource allocation.

Streamlined Data Integration and Sharing

Today's siloed systems make statewide data sharing cumbersome and error prone. A cloud-based architecture with modern integration capabilities will:

- Reduce redundancy by consolidating data into a single source of truth for each data element, while maintaining appropriate system boundaries.
- Simplify data exchange with districts, other state agencies, and federal partners.
- Enable secure API-driven data sharing, reducing reliance on manual SFTP processes.

Stronger Data Governance and Compliance

Modern governance frameworks ensure OSPI's data remains accurate, secure, and compliant with state and federal standards:

- Clear data ownership and stewardship roles prevent ambiguity and errors.
- Automated lineage tracking and quality monitoring reduce audit risk.
- Policies prioritize both access and security, balancing innovation with privacy.
- Real-time data validation reduces human errors.
- Automated reconciliation identifies discrepancies immediately.
- Consistent calculations across reports significantly reduces conflicting metrics.

Increased Operational Efficiency

By reducing manual, spreadsheet-driven processes, OSPI staff can focus on high-value work:

- Automated data pipelines eliminate repetitive tasks and minimize errors.

- Faster reporting cycles reduce delays in compliance submissions and stakeholder updates.
- Staff resources can be redeployed to strategic initiatives rather than maintenance.

Enhanced Analytics and Reporting

Data modernization enables advanced analytics capabilities, supporting a move from compliance-driven oversight to continuous improvement:

- Self-service BI dashboards for districts and internal teams.
- Predictive analytics to identify at-risk students or districts early (e.g., chronic absenteeism patterns, graduation risk indicators).
- KPI-driven insights that connect data to strategic goals and outcomes.
- Faster insights showing impacts on funding for various program outputs (e.g., attendance, participation, equity gaps).

Future-Ready Cloud Infrastructure

Migrating to a scalable, cloud-native architecture positions OSPI for future innovation:

- Elastic scalability supports growth in data volume and new applications.
- Integration with AI/ML tools unlocks predictive and prescriptive analytics.
- Alignment with WaTech standards ensures interoperability and security compliance.

Improved Stakeholder Experience

Districts, educators, and other stakeholders benefit from faster access to reliable information:

- Reduced friction in data requests, improving transparency and trust.
- District-level users gain visibility into statewide benchmarks and trends.
- Supports equity initiatives by ensuring timely, high-quality data for all schools.
- Democratizes data access for smaller programs or teams that do not have technical resources on-staff.
- Anticipated reduction in ad hoc data requests
- Increased district satisfaction (measured via annual survey)

Cost Optimization Over Time

While modernization requires upfront investment, the long-term cost savings are significant:

- Reduced maintenance costs for aging systems.
- Lower operational overhead through automation.
- Avoidance of costly outages and system failures tied to technical debt.

Risk Mitigation and Enhanced Security

Implementation of modern data architecture significantly secures the organization's overall risk portfolio. This empowers OSPI to have:

- Central governance hub allows efficient operational controls and ease-of-use for engineering teams.
- Automated PII detection and masking to protect student privacy continuously.
- Disaster recovery capabilities allow for smooth business continuity transition if disaster occurs.
- Reduced dependency on individual staff knowledge avoids single points of failure for critical processes.
- Cyber threat protection through cloud-native security services.
- Automated audit capabilities support fast and efficient litigation defense and compliance investigations.

13.1.1. Stakeholder Benefits Matrix

OSPI Leadership & Policymakers	OSPI Program & Data Staff	School District Leaders & Administrators	Educators & School Staff	Students & Families
Timely, trusted data for statewide policies	Streamlined data integration across domains	Real-time insights for resource allocation	Actionable student performance data	Earlier identification of learning gaps
Cost optimization and reduced technical debt	Upskilling with familiar Microsoft ecosystem	Faster, accurate reporting to OSPI	Fuller student context through integrated data	Clear, transparent progress information
Clear, consolidated reporting for accountability	Improved data quality and compliance	Statewide benchmarks for equity-focused decisions	Reduced time spent on reporting	More equitable resource distribution

SECTION 14 – RISK MANAGEMENT

The successful implementation and ongoing sustainment of the data modernization effort relies heavily on the proactive identification, evaluation, and management of potential risks throughout its lifecycle. This section outlines the key risks associated with the project, the risk classification, the severity of the risk, and the mitigation strategies that will be employed by the implementation and maintenance teams.

By systematically addressing these risks, we aim to minimize potential disruptions, ensure the solution's effectiveness, and promote long-term sustainability. This will ultimately benefit OSPI and its customers across the State of Washington. Each identified risk will be categorized based on its likelihood and impact, with clear strategies to minimize or manage these risks effectively.

The following risks have been identified and are presented below, organized starting with the highest risk items.

Risk Description	Risk Classification	Severity	Probability	Mitigation Approach
OSPI is unable to source or upskill the talent needed to support platform implementation and long-term sustainment	Workforce	High	High	Sequence key hires from the cost analysis staffing model early; rely on professional services and vendor partners during transition; implement targeted role-based training across teams
Governance policies and standards are applied inconsistently across product teams	Governance	High	High	Centralize policy definition within the Compliance Data Product Team; enforce policies through automation, shared tooling, and team training
Momentum is lost between the end of the feasibility study and funding of implementation	Execution	High	High	Establish and fund a bridge workstream to advance critical path activities, support procurement readiness, and retain key resources
Delivery scope expands without adequate resourcing, resulting in cost overruns or delays	Financial	High	Medium	Maintain strict scope controls through governance and procurement checkpoints; define delivery swim lanes by role and vendor

Risk Description	Risk Classification	Severity	Probability	Mitigation Approach
The selected cloud platform introduces complexity beyond OSPI's ability to support post-implementation	Strategic Alignment	High	Medium	Align platform choice with internal capabilities; avoid over-customization; document platform ownership model and handoff plan
Reporting, compliance, or audit functions are disrupted during system transition	Operational	High	Low	Stage cutover for sensitive systems; conduct full regression testing; run parallel systems as needed during migration period
Stakeholder fatigue results in reduced engagement, slowed adoption, or resistance due to competing initiatives or change overload.	Stakeholder Engagement	Medium	High	Create a clear communication cadence, sequence business changes deliberately across releases using stakeholder impact assessments, and periodically run feedback surveys with stakeholders
Legacy system owners resist decommissioning due to low confidence in new platforms or data products	Change Management	Medium	Medium	Involve SMEs in early migration activities to gather buy-in; allow for temporary parallel operations; build trust by demonstrating data quality and functionality
Collaboration across product teams is limited, preserving silos and delaying integrated delivery	Organizational	Medium	Medium	Establish governance councils and shared backlogs for cross-domain priorities; define clear roles and responsibilities within the product model
High reliance on a small number of technical SMEs creates delivery bottlenecks and knowledge gaps	Operational	Medium	Medium	Document institutional knowledge; embed SMEs into product teams; pair with junior staff for knowledge transfer and succession planning

Risk Description	Risk Classification	Severity	Probability	Mitigation Approach
LEA and partner engagement is inconsistent, leading to integration challenges or unmet needs	Stakeholder Engagement	Medium	Medium	Engage LEAs during discovery and validation; establish feedback loops; assign partner coordination roles within product teams
Adoption of new tools, processes, or team structures is slower than expected	Change Adoption	Medium	Medium	Develop and deliver role-based training; maintain open support channels; monitor adoption metrics and adjust as needed
Architecture standards and reference models are not consistently used by delivery teams	Architecture	Medium	Medium	Establish lightweight Architecture Review Board; promote reusable templates; require intake forms for major design changes