



# WPI

## Visualizing Policy Compliance for Enhanced Cloud Governance

A Major Qualifying Project  
Submitted to the Faculty of  
WORCESTER POLYTECHNIC INSTITUTE  
in partial fulfillment of the requirements for the  
degree of Bachelor of Science

### **Project Team**

Will Huang [whuang3@wpi.edu](mailto:whuang3@wpi.edu)  
Ryan Kornitsky [rtkornitsky@wpi.edu](mailto:rtkornitsky@wpi.edu)  
Ryan Saklad [rjsaklad@wpi.edu](mailto:rjsaklad@wpi.edu)

### **Project Advisors**

Professor Wilson Wong  
Department of Computer Science

Professor Robert Sarnie  
WPI Business School

This report represents the work of WPI undergraduate students submitted to the faculty as evidence of completion of a degree requirement. WPI routinely publishes these reports on its website without editorial or peer review. For more information about the projects program at WPI, please see

<http://www.wpi.edu/academics/ugradstudies/project-learning.html>

## Abstract

In partnership with State Street, we developed user-friendly dashboards to enhance cloud governance and management. Employing Agile Scrum, Power BI, and Azure, we crafted a dashboard that streamlines compliance data management, ensuring regulatory compliance while reducing manual workload for the operations team. Simultaneously, we created an independent React application to visualize historical compliance data, integrating automated reporting and data visualization. This addition demonstrates potential advancements in financial compliance tools, focusing on increasing transparency and client trust. These initiatives have established a new, efficient method for monitoring compliance and governance data within the State Street cloud environment.

## Executive Summary

Our team undertook the challenge of addressing complex cloud governance issues in the financial sector, emphasizing compliance and data security. To enhance visualization tools for efficient cloud governance management, we developed two projects. The first introduced an innovative governance compliance visualization dashboard tailored to State Street's intricate cloud environment. The second, a compliance history dashboard, served as a proof of concept for extending the first dashboard using Azure services not enabled on State Street's cloud environment. These tools streamline compliance data management, improve user accessibility, and offer crucial insights for strategic development and future enhancements. This initiative was driven by the understanding that effective cloud governance is paramount in the financial services industry, where regulatory compliance and data security are not just operational requisites but fundamental to maintaining customer trust and a competitive edge.

Both projects' inception was rooted in research, encompassing an in-depth study of State Street's technology, its pioneering role in the financial sector, and the current landscape of cloud governance. This research underscored the necessity for a tool that not only simplifies the visualization of complex compliance data but also integrates seamlessly with State Street's existing cloud infrastructure, thereby enhancing operational efficiency and regulatory adherence.

We adopted the Agile Scrum methodology for its adaptability in managing complex projects. This approach facilitated a responsive and iterative development process, which is crucial for the dynamic nature of cloud governance challenges. During development, integrating crucial tools like Power BI and Azure services was essential for visualizing compliance data and enabling data storage and retrieval.

A significant challenge faced was the integration and visualization of complex compliance data. By employing iterative feedback loops and continuous testing, we developed robust and user-friendly dashboards. Rigorous assessment confirmed their efficacy in simplifying compliance data visualization, enhancing operational efficiency, and ensuring regulatory compliance. This outcome was particularly significant given the complexity and volume of data typically involved in cloud governance. Our evaluation also highlighted the dashboard's impact in enhancing operational efficiency and regulatory compliance for State Street, demonstrating its potential to significantly contribute to more effective and streamlined cloud governance practices.

The projects' implications extend beyond State Street, serving as a model for other financial institutions facing similar cloud governance challenges. The focus on user-centric design and iterative development provides insights into crafting technology solutions for compliance and governance teams. As cloud technologies evolve, the lessons learned and solutions developed in this project will guide future endeavors in cloud governance. As cloud technologies continue to evolve and play a more central role in the operational strategies of financial institutions, the lessons learned and solutions developed in this project will remain pertinent, guiding future endeavors in cloud governance.

These projects not only met State Street's immediate needs for effective cloud governance but also contributed to the broader discourse on cloud migration and resilience in the financial sector. They exemplify the potential of collaborative, research-driven approaches in addressing complex technological challenges. The development process sets a precedent for future initiatives in enhancing cloud governance, marking a significant step forward in achieving resilient, efficient, and compliant practices for State Street and the broader financial services industry.

# Acknowledgments

This project owes its success to the invaluable support and guidance extended by numerous individuals and organizations. Foremost, we express our gratitude to our advisors, Professor Wong and Professor Sarnie, whose essential feedback and unwavering guidance significantly influenced the project's outcome. Their consistent support during our weekly meetings played a pivotal role in our success. Special thanks are extended to Marc Trudeau for facilitating agile learning activities and Michael O'Connor for providing project advice.

A heartfelt acknowledgement goes to State Street Corporation for their enthusiasm and willingness to collaborate with WPI and our project group on this initiative. We extend our sincere thanks to Prafull Srivastava, our primary sponsor, for their passion and efforts in championing this project. Their role in connecting us with key individuals within the company and offering substantial assistance with project planning is truly appreciated. We also express our thanks to Akshay Chawade for providing insightful advice and assistance throughout the project.

We extend our heartfelt thanks to the dedicated teams at State Street for their exceptional support and guidance throughout our project. Our gratitude particularly goes to the Azure Platform team, led by Akshay Sharma and Ravi Palanki, and the Platform Service Engineering team, under the leadership of Akshay Chawade, for their invaluable assistance. We also deeply appreciate the efforts of the Platform Operations team, led by Anurag Phadke, and the Platform Landing Zone Engineering team, directed by Purush Rudrakshala. The contributions of the Service and Policy Engineering team, including Sudheer K Yerrabati and Amit Kumar, were critical in our progress. Our special thanks to the Compliance Ops team, comprising Zhongyuan Ying (Trista), Priyanka

Karli, and Rajesh Tripathi, for their enthusiastic cooperation and insight. We are especially grateful to Kimberly Wills for her diligent work and guidance throughout the project's journey.

We also want to recognize the developers who contributed to the development of Azure and Power BI, as well as contributors to the CCO Insights open-source repository. Without these tools, our work would not have been possible. Special thanks to all instructors, contributors, and developers who created extensive documentation and tutorials about Azure services, Power BI, KQL, Power Query M, and CCO Insights. Without these resources, acquiring the necessary concepts to complete this project within the given timeframe would not have been achievable.

Thank you.

# Table of Contents

<b>Abstract.....</b>	<b>ii</b>
<b>Executive Summary.....</b>	<b>iii</b>
<b>Acknowledgments.....</b>	<b>v</b>
<b>Table of Contents.....</b>	<b>vii</b>
<b>List of Figures.....</b>	<b>ix</b>
<b>List of Tables.....</b>	<b>x</b>
<b>Authorship.....</b>	<b>xi</b>
<b>1. Introduction.....</b>	<b>1</b>
Problem.....	2
Goal.....	2
<b>2. Research.....</b>	<b>4</b>
2.1 Company Background.....	4
2.2 Background Research.....	6
<b>3. Methodology.....</b>	<b>15</b>
<b>4. Software Development Environment.....</b>	<b>19</b>
4.1 Azure Automation.....	19
4.2 Azure Policy.....	20
4.3 Azure Portal.....	21
4.4 Azure Resource Graph.....	22
4.5 Azure SQL Managed Instance.....	23
4.6 Citrix Virtual Desktop.....	25
4.7 Continuous Cloud Optimization Insights.....	25
4.8 Jira.....	27
4.9 Kusto Query Language.....	27
4.10 Microsoft Excel.....	29
4.11 PowerBI.....	30
4.12 Azure Blob Storage.....	31
4.13 Azure Virtual Machine.....	32
4.14 Express.....	33
4.15 React.....	35
4.16 Recharts.....	36
4.17 Vite.....	37
4.18 Communication Technologies - Discord, iOS Messaging, Microsoft Outlook, Microsoft Teams, Slack, Zoom.....	38
<b>5. Software Requirements.....</b>	<b>39</b>

5.1 Software Requirements Gathering Strategy.....	39
5.2 Functional and Non-functional Requirements.....	40
5.3 User Stories and Epics.....	41
<b>6. Design.....</b>	<b>46</b>
6.1 Governance Compliance Dashboard.....	46
6.2 Compliance History Dashboard.....	48
6.3 Data Model Entity Relationship Diagram.....	49
<b>7. Software Development.....</b>	<b>52</b>
7.1 Pre-sprint Setup and Configuration.....	55
7.2 Sprint One.....	56
7.3 Sprint Two.....	60
7.4 Sprint Three.....	64
<b>8. Assessment.....</b>	<b>68</b>
8.1 Business Learnings.....	68
8.2 Technical Learnings.....	70
8.3 Accomplishments.....	74
8.4 Limitations.....	78
<b>9. Business and Risk Management.....</b>	<b>80</b>
9.1 Risks and Costs.....	80
9.2 Benefits.....	84
9.3 Compliance History Dashboard.....	90
<b>10. Future Work.....</b>	<b>93</b>
10.1 Integration of Additional Azure Features.....	93
10.2 Integration of the Second Dashboard.....	93
10.3 Machine Learning for Predictive Analytics.....	94
10.4 Enhanced User Interactivity and Personalization.....	94
10.5 Integration with External Risk Management Systems.....	94
10.6 Automated Remediation.....	95
<b>11. Conclusion.....</b>	<b>96</b>
<b>References.....</b>	<b>98</b>
<b>Appendix A: Glossary of Terms.....</b>	<b>103</b>



## List of Figures

Figure name	Page number
Figure 2.2.5: Example of a Dashboard	12
Figure 6.1: Project 1 High-Level Project Architecture	46
Figure 6.2: Project 2 High-Level Project Architecture	48
Figure 6.3: Entity Relationship Diagram	49
Figure 7.1: Weekly Overview	54
Figure 8.3.1: Compliance History Dashboard	75
Figure 8.3.2: Compliance History Dashboard PDF Display	76

## List of Tables

Table name	Page number
Table 5.3: User Stories and Epics	41
Table 7.2: Sprint 1 Backlog	56
Table 7.3: Sprint 2 Backlog	60
Table 7.4: Sprint 3 Backlog	64

## Authorship

Section	Main Author(s)	Main Editor(s)
Cover Page	Ryan Saklad	Will Huang, Ryan Kornitsky, Ryan Saklad
Abstract	Ryan Saklad	Will Huang, Ryan Kornitsky, Ryan Saklad
Executive Summary	Ryan Saklad	Will Huang, Ryan Kornitsky, Ryan Saklad
Acknowledgements	Ryan Kornitsky	Will Huang, Ryan Kornitsky, Ryan Saklad
1.0 Introduction	Ryan Saklad	Will Huang, Ryan Kornitsky, Ryan Saklad
2.0 Research	Ryan Saklad, Will Huang	Will Huang, Ryan Kornitsky, Ryan Saklad
3.0 Software Development Methodology	Ryan Saklad	Will Huang, Ryan Kornitsky, Ryan Saklad
4.0 Software Development Environment	Ryan Saklad	Will Huang, Ryan Kornitsky, Ryan Saklad

5.0 Software Requirements	Ryan Saklad, Will Huang, Ryan Kornitsky	Will Huang, Ryan Kornitsky, Ryan Saklad
6.0 Design	Ryan Kornitsky	Will Huang, Ryan Kornitsky, Ryan Saklad
7.0 Software Development	Ryan Saklad, Will Huang	Will Huang, Ryan Kornitsky, Ryan Saklad
8.0 Assessment	Ryan Kornitsky	Will Huang, Ryan Kornitsky, Ryan Saklad
9.0 Business and Risk Management	Ryan Saklad, Ryan Kornitsky	Will Huang, Ryan Kornitsky, Ryan Saklad
10.0 Future Work	Ryan Kornitsky	Will Huang, Ryan Kornitsky, Ryan Saklad
11.0 Conclusion	Ryan Saklad	Will Huang, Ryan Kornitsky, Ryan Saklad
Table of Contents	Ryan Saklad, Ryan Kornitsky	Will Huang, Ryan Kornitsky, Ryan Saklad

# 1. Introduction

In the fast-evolving landscape of financial services, where cloud computing is revolutionizing data management and processing, financial firms are increasingly facing complex challenges in governance. These firms are required to adhere to a myriad of policies and compliance regulations, making the management of compliance and data security a critical priority. In response to these industry-wide challenges, our team, employing Agile Scrum methodology, partnered with State Street, a prominent player in the financial sector. We embarked on developing a specialized compliance visualization dashboard, designed to integrate with cloud infrastructures commonly used by financial institutions. This tool is crucial in aiding these firms to manage and monitor their adherence to regulatory standards in the cloud environment.

The first project, also known as the Governance Compliance Dashboard, leveraged Power BI and Azure services to create a user-friendly and efficient dashboard, aimed at streamlining the management of complex compliance data. It provided an intuitive interface for monitoring compliance status, facilitating quick decision-making, and effective policy enforcement. The use of Azure services ensured a scalable and secure platform for handling sensitive financial data, while Power BI tools brought in advanced analytics capabilities for deeper insights into compliance trends and potential risks.

Concurrently, alongside the primary dashboard project, our team initiated a separate but complementary second project: the development of the Compliance History Dashboard. This independent React application, deployed on an Azure virtual machine for robust and scalable hosting, is dedicated to visualizing historical compliance data. Serving as a proof of concept for extending the capabilities of the first dashboard, the Compliance History Dashboard employs

Recharts for dynamic visualizations. It features a compliance history bar chart, a severity bar chart, and a line chart tracking compliance trends over time. The deployment on an Azure VM ensures reliability, security, and scalability, crucial for handling the complex data demands of this application. The Compliance History Dashboard is designed to provide State Street with a comprehensive overview of its compliance history, aiding in strategic decision-making and offering advanced tools for auditing, managing, and record-keeping. This initiative is a testament to our commitment to leveraging state-of-the-art cloud technologies to address the intricate requirements of compliance management in the financial sector.

## **Problem**

The core challenge confronting State Street in this digital transformation era is the effective monitoring and management of complex compliance requirements within cloud environments. The current state of cloud governance is marked by an overwhelming influx of compliance data, making it arduous for financial institutions to not only gather but also decipher and utilize this information effectively. Existing methods and tools often prove inadequate in providing a coherent and actionable understanding of compliance statuses, leading to potential governance and security risks. The projects seek to bridge this gap, offering a solution that elevates State Street's capabilities in interpreting and managing compliance data efficiently.

## **Goal**

Our primary goal with the Compliance Governance Dashboard was to develop an intuitive, user-friendly interface that simplifies and enhances the visualization of compliance data. This dashboard is designed to be highly scalable, managing the complexity and volume of State Street's

cloud governance operations, and adaptable to the evolving landscape of cloud governance and compliance requirements.

Additionally, the Compliance History Dashboard, with its automated generation of weekly compliance report PDFs through Azure Blob Storage, provides State Street with valuable tools for various purposes including auditing and management. Both projects demonstrate our team's commitment to technical innovation and deep understanding of regulatory requirements in the financial sector, advancing compliance management solutions.

## 2. Research

### 2.1 Company Background

State Street Corporation, founded in 1792 as Union Bank and later becoming National Union Bank of Boston in 1865, represents a significant force in American finance. Chartered by Massachusetts Governor John Hancock, it began operations in Boston's thriving maritime capital. Over its extensive history, State Street underwent several transformations, including mergers with Second National Bank in 1955 and Rockland-Atlas National Bank in 1961, leading to its current name that pays homage to Boston's maritime legacy. Today, it stands as a Fortune 500 company and the second oldest continuously operating bank in the United States, boasting about 39,000 employees and ranking among America's largest banks by assets. Its evolution into a leading financial services provider has been marked by a focus on serving institutional investors through key divisions such as State Street Global Advisors, Global Services, Global Markets, and Global Exchange. The Global Services division manages a diverse portfolio, including stocks, derivatives, ETFs, fixed income, private equity, and real estate, while the Global Markets division provides services in research, trading, and securities lending across various asset classes (History, 2023).

Structured around three core pillars - Investment Servicing, Investment Management, and Investment Research - State Street has notably expanded its services, incorporating foreign exchange, performance measurement, and compliance monitoring into its Investment Servicing segment. The Investment Management division employs advanced quantitative models for decision-making across different asset classes (History, 2023). In the technological realm, State Street has been pioneering, announcing strategic partnerships with AWS and Microsoft Azure for a



multi-year technology transformation journey. This journey aims at modernizing IT infrastructure and consolidating data centers into a more efficient footprint with both private and public cloud capabilities. The company's commitment extends to integrating data analytics, artificial intelligence, and blockchain to revolutionize real-time insights and risk assessments (State Street, 2023).

Recognized as one of the 11 U.S. and 30 worldwide systemically important financial institutions (SIFIs) by the Financial Stability Board, State Street holds a crucial role in the global financial ecosystem. This distinction adds significant responsibility to its operations, setting it apart from competitors like J.P. Morgan and Bank of New York Mellon. Its Investment Research arm leverages artificial intelligence to produce a broad spectrum of research that informs internal strategies and influences industry norms (Liberto, 2023; History, 2023).

State Street is a seasoned expert in building financial infrastructure, servicing nearly 10% of the world's assets. It offers asset managers, insurance companies, official institutions, and central banks a unique market perspective, aiding them to stay ahead of change and drive business performance. The company has been instrumental in creating the first open ecosystem to connect the investment lifecycle, exploring tokenization, and providing electronic trading platforms. It has also significantly contributed to the growth of the ETF industry, now valued at \$6 trillion (State Street, 2023).

In the late 20th century, State Street strategically shifted focus by divesting its retail and commercial banking businesses to concentrate on securities processing and asset management. This shift led to the acquisition of Charles River Development in 2018, a significant investment management software provider. The launch of its Alpha data platform in 2020, a cloud-based investment servicing software, further established State Street as a leader in digital assets (Lopez et al., 2022).

Embarking on a multi-year technology transformation, State Street has partnered with Amazon Web Services (AWS) and Microsoft for cloud and infrastructure solutions. This transformation is aimed at consolidating the company's data center facilities into a more optimized and purpose-fit design. The hybrid and hyper-converged model combines traditional data center elements, allowing hosting workloads across private or public environments, thus enabling accelerated application migration and modernization. This cloud-first strategy is viewed not only as a technological advancement but also as a financial imperative, aiming to increase resiliency, foster rapid business product innovation, and enhance customer experiences. The cloud architectures from AWS and Microsoft Azure will enable clients to manage and access data in real-time, streamlining their investment processes and accelerating decision-making (State Street, 2023).

## **2.2 Background Research**

### **2.2.1 What is Cloud Computing**

Cloud computing is a significant shift in technology utilization, characterized by the remote delivery of a range of services. These services include server space, data storage, database management, networking capabilities, software applications, analytics, and artificial intelligence, all accessed over the Internet. This approach offers the benefits of rapid deployment, innovative capabilities, flexible resource allocation, and cost efficiency due to economies of scale (Microsoft, 2023b).

At its core, cloud computing moves away from the traditional model of local servers or personal devices handling workloads. Instead, it embodies an on-demand provision of IT resources via the Internet, often with pay-as-you-go pricing. This approach reduces the need for upfront capital expenditure and allows for dynamic resource adjustment in response to changing needs.

Beyond traditional data storage, cloud computing has become a catalyst for business innovation, providing tools to build new customer experiences and optimize operations with advanced technologies like machine learning and natural language processing (Celner, 2019). The versatility of cloud computing is evident in its service models, which include:

- **Infrastructure as a Service (IaaS):** This model offers virtualized computing resources over the Internet. It provides a foundational IT infrastructure upon which companies can build their digital platforms (Microsoft, 2023b).
- **Platform as a Service (PaaS):** PaaS facilitates the development, running, and management of applications without the complexity of building and maintaining the underlying infrastructure. This streamlines development operations, making it easier for organizations to focus on innovation rather than infrastructure management (Amazon Web Services, 2023).
- **Software as a Service (SaaS):** SaaS delivers software applications over the internet on a subscription basis. This model enables organizations to access and use cutting-edge applications without the overhead of maintenance, updates, and patching, thus simplifying software management (IBM, 2023).

Cloud computing is not just a technology but a paradigm shift that offers scalable, flexible, and cost-effective solutions for modern computing needs. It enables organizations to leverage the latest technological advancements, ensuring agility and competitiveness in a rapidly evolving digital landscape.

## 2.2.2 Cloud Computing in Financial Technology

The integration of cloud computing within the financial technology sector is propelling the industry forward at an unprecedented pace. Fintech companies are harnessing the power of the cloud to meet and exceed the stringent requirements of today's financial landscape, including the imperative for security, service quality, innovation, and scalability (EidosMedia, 2022).

Cloud computing is paramount for fintech companies like State Street for several reasons. Firstly, it has revolutionized the delivery of digital services, allowing banks to scale rapidly, especially in response to the accelerated digital shift caused by the COVID-19 pandemic. This is evident as major financial institutions like Wells Fargo, Morgan Stanley, and Capital One transition to cloud services, with NASDAQ moving its markets to Amazon Web Services and Goldman Sachs collaborating to launch Financial Cloud for Data (EidosMedia, 2022). Data serves as the lifeblood of the fintech industry, supporting a vast array of activities essential to modern finance. Cloud computing offers fintech firms the capabilities to securely manage significant volumes of data, enabling a plethora of operations from transaction processing to complex analytics, all accessible anytime and anywhere (McIntosh, 2021).

Fintechs are adopting cloud computing to improve agility, scalability, and customer service, with 72% of banking sector IT executives reporting that cloud computing helps achieve their business priorities (EidosMedia, 2022). Cloud infrastructure offers enhanced security through advanced data encryption and zero-trust security measures, allowing financial institutions to securely store and instantly access large volumes of data (EidosMedia, 2022). Additionally, cloud computing enables innovative services like open banking and Banking as a Service, which are reshaping customer relationships with financial service providers (EidosMedia, 2022). In exploring the transformative impact of cloud computing on the fintech sector, key areas of focus are:

- **Accelerated Innovation:** The agility afforded by cloud computing is a catalyst for rapid innovation within fintech. Financial organizations can develop and deploy new products swiftly, adapting to market changes and emerging trends with ease. The resilience of cloud infrastructure was especially evident during the COVID-19 pandemic, where it played a critical role in helping financial services navigate and surmount operational challenges (McIntosh, 2021).
- **Enhanced Security:** With the rise in data breaches and cyber threats, the cloud enhances how financial businesses protect customer data. Innovations in data encryption, zero-trust verification, and access control mean that cloud computing can offer more robust security measures than traditional IT infrastructures, thereby fortifying the trust placed by customers in their financial service providers (McIntosh, 2021).
- **Greater Scalability:** Fintech is characterized by rapid growth, requiring infrastructures that can keep pace. Cloud infrastructure empowers financial companies to scale operations seamlessly, facilitating the expansion of customer bases and the digitization of traditional banking services in a cost-effective manner (McIntosh, 2021).

As the fintech sector continues to leverage the myriad benefits of cloud computing, the focus inevitably shifts towards the pivotal aspects of cloud compliance and governance. In the context of financial firms like State Street, where data sensitivity and regulatory adherence are paramount, understanding and implementing robust cloud compliance and governance frameworks is crucial. This next section will delve into what cloud compliance and governance entail, exploring their significance in ensuring that financial institutions not only comply with stringent regulatory standards but also maintain the highest levels of data integrity and security. This is especially critical

in an industry where trust and compliance are not just operational necessities but foundational to business integrity and customer confidence.

### 2.2.3 Cloud Compliance

Cloud compliance refers to ensuring that services offered by cloud computing platforms adhere to a range of regulations and standards set by state and federal governments, regulatory bodies, or other jurisdictional authorities, in addition to any internal policies. It is a critical aspect of cloud computing, as non-compliance can lead to regulatory fines, lawsuits, cybersecurity incidents, and reputational damage. Enterprises must understand the details of cloud provider services and how well they meet the organization's requirements (Shacklett, 2023).

The importance of cloud compliance cannot be overstated. Failure in compliance measures can be costly, with the average cost of a data breach reaching \$4.45 million in 2023. Beyond financial losses, non-compliance can lead to a loss of trust among enterprise users, executives, stakeholders, and customers, negatively impacting revenues and operational performance. Compliance is not only a technical issue but also involves decision-making, monitoring, audits, governance, security, data protection, risk management, and legal guidance. It is a shared responsibility between the enterprise and the cloud service provider (Shacklett, 2023).

In financial firms, cloud compliance is particularly crucial due to the sensitive nature of financial data and the stringent regulatory environment. Financial organizations must conform to specific data and reporting requirements, such as those outlined in the Sarbanes-Oxley Act (SOX). Compliance in the financial sector involves not only adhering to these regulations but also ensuring that cloud service providers meet the same standards. This is vital for protecting customer data, maintaining trust, and avoiding legal and financial repercussions. The integration of cloud

governance in this context becomes essential, as it provides a framework for managing and ensuring compliance across all cloud services and operations (Shacklett, 2023).

#### **2.2.4 Cloud Governance**

Cloud governance refers to the systematic oversight of cloud-based technology offerings. This oversight involves formulating and implementing policies, procedures, and technical measures to effectively administer and regulate the operational environment of cloud services. The objective is to ensure that these services are used in a way that aligns with an organization's strategic IT objectives and business goals. This governance ensures that cloud services align with an organization's IT strategy and business objectives, providing a structure for decision-making regarding cloud usage. It is about obtaining complete visibility into all cloud activities, optimizing performance, lowering operational costs, and minimizing security risks (SailPoint Technologies, Inc., 2021).

The importance of cloud governance lies in its ability to enhance visibility and control over cloud assets, improve operational efficiency, and ensure that cloud infrastructures are secure. It helps organizations comply with government regulations such as SOX, HIPAA, GDPR, and others, making it easier to meet legal and ethical standards. Additionally, effective cloud governance can reduce labor and operational costs by automating time-consuming manual processes, thereby boosting innovation and profits. It addresses challenges such as managing cloud spending, securing cloud infrastructure, and controlling unauthorized access to critical assets (SailPoint Technologies, Inc., 2021).

In financial firms, cloud governance is particularly crucial due to the highly regulated nature of the financial industry and the sensitivity of financial data. These firms are subject to stringent regulations that dictate how data should be handled and protected. Cloud governance ensures that

financial institutions meet these regulations when they store, process, or transmit data in the cloud. It is vital for maintaining customer trust, preventing data breaches, and avoiding legal penalties. Moreover, cloud governance in financial firms helps in managing cloud spending, securing cloud infrastructure, and controlling unauthorized access to critical assets, which are essential for maintaining the integrity and confidentiality of financial data (SailPoint Technologies, Inc., 2021).

## 2.2.5 Governance Compliance Dashboard

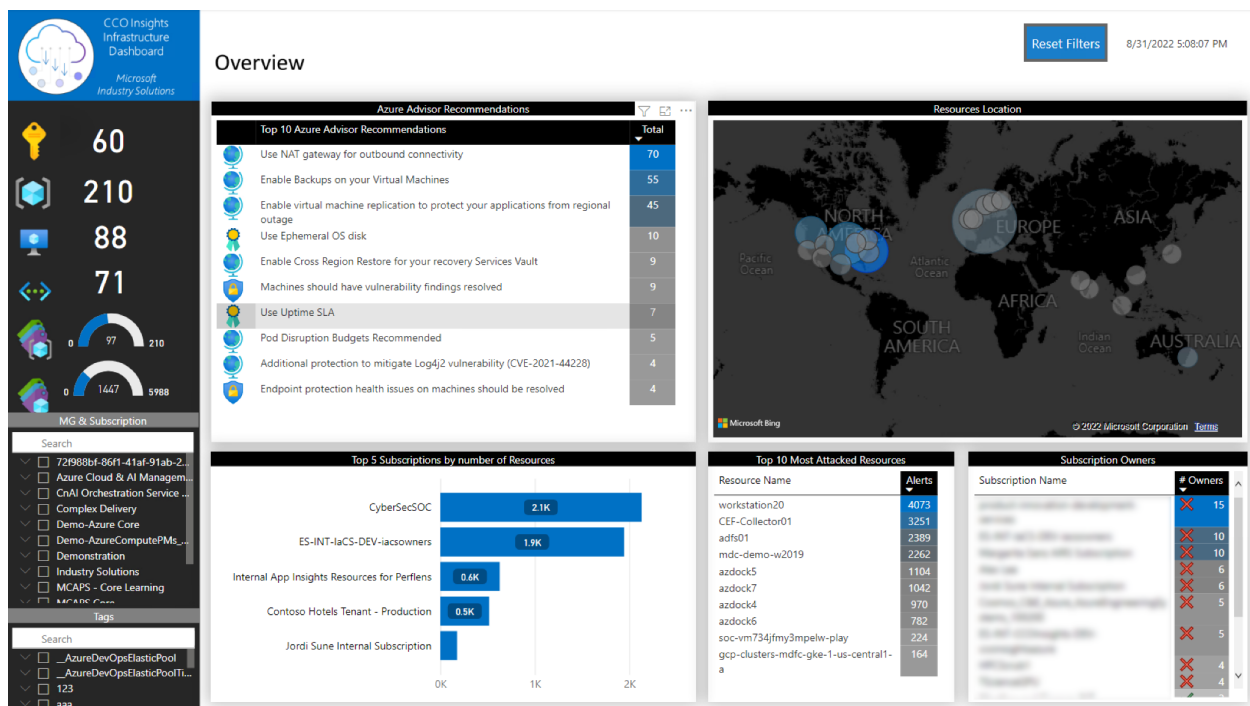


Figure 2.2.5: Example of a Dashboard (Suñé, J. 2023)

The primary purpose of a Governance Compliance Dashboard is to offer a comprehensive overview of resources and compliance policies, highlighting areas of non-compliance in a more efficient manner than traditional portal-based methods. This dashboard typically includes various filters, allowing users to sort and view data by subscription, management groups (MG), policy assignment groups, severity, and other relevant criteria. Such a dashboard is instrumental in



simplifying the monitoring and management of compliance across an organization's cloud environment.

State Street leverages the Governance Compliance Dashboard to enhance visibility and control over its compliance status, offering significant benefits to both management and application teams. For management, the dashboard provides a quick, comprehensive overview of the compliance status across all resources or can be focused on specific management groups or subscriptions, enabling easy assessment and addressing of compliance issues organization-wide. Meanwhile, application and development teams utilize the dashboard to monitor the compliance level of their specific subscriptions, ensuring that the resources they provision are in line with their policies. This tool offers a glance at compliance status without the need to delve into detailed portal analytics, streamlining the compliance management process for both management and application teams.

The introduction of the new Governance Compliance Dashboard has marked a significant advancement in the way compliance monitoring and management are handled. Firstly, the dashboard facilitates enhanced compliance tracking by providing up-to-date information on compliance status, which is updated weekly. This feature is instrumental in the timely identification and resolution of compliance-related issues. Furthermore, the dashboard is equipped with advanced analytics and reporting features, offering deeper insights into compliance data, thereby aiding in more informed and strategic decision-making. Additionally, significant attention has been given to user experience in the dashboard's design. Its user-friendly interface improvements make it easier for various teams within the organization to navigate and efficiently extract the necessary information. These combined features of the Governance Compliance Dashboard not only streamline compliance processes at State Street but also set a new standard in compliance management tools.

Recognizing the diverse needs of various user roles within the organization, the Governance Compliance Dashboard is equipped with customization options. These options allow users to tailor the dashboard to their specific requirements, ensuring that each team or individual has the most relevant and useful information at their fingertips. This level of customization is crucial in a complex environment like State Street, where different teams may have varying compliance monitoring needs.

The Governance Compliance Dashboard is a pivotal tool in State Street's cloud governance strategy. It not only simplifies compliance monitoring but also enhances the overall efficiency and effectiveness of the organization's cloud governance efforts. This dashboard serves as a model for other organizations looking to improve their cloud compliance and governance processes.

### **2.2.6 Compliance History Dashboard**

The addition of the section project, the Compliance History Dashboard, also stems from a lot of the same principles as the Governance Compliance Dashboard, since it is aimed to be an extension of that dashboard. The goal of this project was to explore potential meaningful extensions to the first dashboard using Azure features that were not enabled on State Street's cloud environment. Due to this, we had to develop the second dashboard outside of the State Street environment. This dashboard showcases compliance history visualizations that can be used to see how policy and resource compliance changes over time. Additionally, it has automated compliance reporting that gives an overview of the policy compliance and severity over time that is uploaded to Azure Blob Storage for secure access. This report can be used for auditing, recording keeping, and much more. Since these Azure services are in the process of being enabled for State Street, this extension to the first dashboard will be very useful and serve as great documentation for further expansion on this dashboard in the future.

### 3. Methodology

The Agile Scrum methodology, extensively employed in software development projects, is specifically designed for complex environments where adaptability, responsiveness, and iterative progress are paramount. As defined in the "Scrum Reference Card," Scrum is a framework for product development that utilizes cross-functional teams. This approach places a strong emphasis on empirical feedback based on real-world data and team self-management. This methodology stands in stark contrast to the traditional 'waterfall' development approach, offering a more flexible and responsive method. It is particularly well-suited for complex domains where the ability to probe, sense, and respond is essential (Rubin, 2013, p. 8).

Within the Agile Scrum framework, three primary roles are crucial: the Product Owner, the Scrum Master, and the Development Team. The Product Owner plays a central role in product leadership, responsible for maximizing the value of the development effort. This includes continuously re-prioritizing the Product Backlog, a force-ranked list of desired functionality visible to all stakeholders, and acting as the final arbiter of requirements questions (James, 2023). The Scrum Master's role is multifaceted, working within the organization to facilitate the adoption of Scrum, ensuring it is properly understood and enacted. They also create an environment conducive to team self-organization and help everyone involved understand and embrace the Scrum values, principles, and practices, effectively acting as a coach (James, 2023). These Scrum values, which include commitment, courage, focus, openness, and respect, form the foundation for collaborative teamwork and guide the behavior of team members toward achieving the goals of each sprint and fostering continuous improvement. The Development Team, described as cross-functional and

self-organizing, is responsible for developing the increment and handling both internal and external collaboration, determining the best way to accomplish the goals set out by the Product Owner (Rubin, 2013, p. 198).

The Agile Scrum process is structured around sprints, which are fixed periods designed to achieve specific goals. According to "Essential Scrum," sprints always have a fixed start and end date and generally should all be of the same duration. The consistency in sprint duration helps establish a predictable rhythm for the team and stakeholders (Rubin, 2013, p. 20). Sprint Planning is a critical phase where the team and Product Owner decide which Product Backlog Items to convert into working product increments. The Daily Scrum, characterized as a daily 15-minute meeting, ensures the team remains aligned on their goals and can quickly address any issues that arise (James, 2023). The Sprint Review is designed to inspect the Product Increment and adapt plans for it, often involving customers, end-users, and other stakeholders (James, 2023). Each Sprint concludes with a Retrospective, a meeting for the team, Product Owner, and Scrum Master to reflect on their way of working together and adapt it for future Sprints (Rubin, 2013, p. 27).

In the Agile Scrum framework, the concepts of user stories and epics play a pivotal role in organizing and managing the Product Backlog. User stories are a common format for Product Backlog items, providing a simple, concise, and customer-centric description of a feature from an end-user perspective. They are instrumental in ensuring that the development work aligns closely with customer needs and priorities. Epics, on the other hand, are oversized Product Backlog items that often encompass a broader scope or a more complex feature, requiring breakdown into smaller, more manageable user stories. This approach facilitates a more flexible and adaptive planning process, contrasting with traditional methods that often break down features into sequential tasks,

lacking the ability to be independently prioritized and potentially missing the core business value from the customer's perspective (James, 2023).

Agile Scrum was selected for this project due to its ability to deliver small features first and allow for ongoing discovery and plan revision, crucial in software development where requirements and environments can rapidly evolve. The core values of Agile Scrum, as highlighted in the "Manifesto for Agile Software Development," include prioritizing individuals and interactions over processes and tools, working software over comprehensive documentation, customer collaboration over contract negotiation, and responding to change over following a plan (Beck et al., 2013). These values emphasize the adaptability and responsiveness of Agile Scrum, making it an ideal choice for projects in dynamic and complex environments.

In addition to the three primary roles, Scrum also involves several key artifacts and ceremonies. The main artifacts include the Product Backlog, the Sprint Backlog, and the Increment. The Product Backlog is a dynamic, ordered list of everything that might be needed in the product, and is the single source of requirements for any changes to be made to the product (Rubin, 2013, p. 18). The Sprint Backlog is a set of items selected from the Product Backlog to be completed during the Sprint, along with a plan for delivering them (Rubin, 2013, p. 22). The Increment is the sum of all the Product Backlog items completed during a Sprint and all previous Sprints, representing the next version of a product (Rubin, 2013, p. 349). Further key ceremonies in Scrum, besides the Sprint Planning, Daily Scrum, Sprint Review, and Sprint Retrospective, include Backlog Refinement and the Sprint Goal. Backlog Refinement, sometimes referred to as Backlog Grooming, involves reviewing the Product Backlog items to ensure they are appropriately prepared and prioritized for future Sprints (James, 2023). This process involves breaking down large items, detailing them, and estimating the effort required. The Sprint Goal is a short, one- or two-sentence description of what

the team plans to achieve during the Sprint. It provides guidance to the Development Team on why it is building the Increment (James, 2023).

Another aspect of Scrum is its emphasis on continuous improvement. The Sprint Retrospective, in particular, focuses on reflecting on the past Sprint to identify and implement improvements to be enacted during the next Sprint (Rubin, 2013, p. 27). This meeting is not just about identifying what went wrong, but also what went right and how the team can continue to improve its work processes and environment. In Scrum, the concept of "Definition of Done" is also critical. It is a clear and concise list of criteria that an increment must meet to be considered complete. This definition ensures transparency and provides a shared understanding of what it means for work to be complete, increasing the quality of the increment (Rubin, 2013, p. 74). Scrum also focuses on the principle of transparency. All aspects of the Scrum process must be visible to those responsible for the outcome. This transparency requires that the artifacts are defined in a way that everyone has a common understanding of what is being seen (Rubin, 2013, p. 35). For example, for a Product Backlog item or an Increment, everyone must understand what it represents.

Agile Scrum is a flexible, iterative, and effective framework for managing complex software development projects. It emphasizes collaboration, continuous improvement, and a focus on delivering high-quality software increments regularly. By embracing the Scrum values, teams can navigate the challenges of modern software development, ensuring that their products meet the evolving needs of customers and stakeholders.

## 4. Software Development Environment

### 4.1 Azure Automation

Azure Automation stands as a key cloud-based automation service, adept at consistent management across Azure and non-Azure environments (Sudhir, 2022). It integrates process automation, configuration management, update management, and shared capabilities, streamlining deployment, management, and decommissioning of enterprise workloads and resources. This service significantly reduces operational costs by automating repetitive tasks and is enhanced by functionalities like Hybrid Runbook Worker for direct execution on various environments and PowerShell DSC for configuration management. In our project, Azure Automation was instrumental in automating data retrieval from Azure Resource Graph Explorer. The service efficiently extracted compliance and configuration data, which was then seamlessly uploaded to our SQL Managed Instance (SQL MI), ensuring real-time data synchronization and management. This automation not only optimized our workflows but also ensured accuracy and consistency in data handling.

The service's shared capabilities, including role-based access control, flexible scheduling, and source control integration, further elevate its utility. Its heterogeneous support extends to both Windows and Linux environments, enabling consistent management across different platforms. Azure Automation's integration with Azure services like Arc, Alerts, and Monitor, caters to a broad spectrum of use cases, from task scheduling and resource deployment to monitoring and governance automation. The cost-effective pricing model, including free run time minutes, makes Azure Automation a versatile and economical choice for businesses seeking efficient cloud resource management (Sudhir, 2022).

## 4.2 Azure Policy

Azure Policy, a service within Azure, plays a crucial role in enforcing organizational standards and assessing compliance on a large scale (Davidsmatlak, 2023a). It achieves this through its compliance dashboard, which offers an aggregated view of the environment's overall state. This dashboard allows users to drill down to a detailed level, evaluating compliance on a per-resource and per-policy basis. The ability to assess compliance at this granularity is essential for maintaining control and visibility over a complex cloud environment. The common use cases for Azure Policy span various critical areas such as resource consistency, regulatory compliance, security, cost management, and general management. These use cases are addressed through policy definitions, which are essentially business rules written in JSON format. Azure Policy evaluates resources and their actions by comparing them against these business rules. The service simplifies management by allowing the grouping of several business rules, aiding in the coherent governance of cloud resources. Azure Policy's approach involves the application of specific business responses to compliance evaluation outcomes. This flexibility allows organizations to tailor their response strategies to different compliance scenarios, ensuring that their cloud environment aligns with their specific governance and compliance requirements (Davidsmatlak, 2023a).

In our project, Azure Policy was the source of compliance data for our dashboard. The service's detailed compliance information was integral to our project's success. We leveraged Azure Policy to understand the compliance status of our resources and ensure they adhered to the set standards and policies. The insights gained from Azure Policy were critical in shaping our dashboard, enabling us to present a comprehensive view of our project's compliance posture. By integrating Azure Policy into our workflow, we could maintain a high standard of resource management and governance, which was pivotal in the project's overall success.



## 4.3 Azure Portal

The Azure portal, a web-based, unified console developed by Microsoft, is an essential tool for managing Azure subscriptions and services (JnHs et al., 2022). It provides a versatile and user-friendly graphical interface, enabling users to build, manage, and monitor a wide range of applications and services, from simple web apps to complex cloud deployments. The Azure portal is particularly notable for its resilience and continuous availability, with a presence in every Azure data center. This global distribution ensures the portal's resilience against individual data center failures and mitigates network slow-downs by being geographically proximate to users. Additionally, the Azure portal is designed for continuous updates without requiring downtime for maintenance, ensuring a seamless and uninterrupted user experience.

The Azure portal's layout is designed for efficiency and ease of use. The portal menu is a key element of the interface, providing rapid access to essential functionalities and resource types. Users have the option to set the portal menu to either a flyout or docked mode, depending on their preference. In flyout mode, the menu remains hidden until needed, whereas in docked mode, it is always visible, allowing for quick access to various services and resources. Dashboards in the Azure portal offer a focused view of the resources crucial to the user's subscription. Users are initially provided with a default dashboard, which can be customized to consolidate frequently used resources into a single view. This customization does not affect other users, but one can create additional dashboards for personal use or share them with others in the organization. This feature enhances the user experience by enabling a more organized and efficient management of resources. Furthermore, the Azure portal's user interface is built with global elements like the Azure portal menu and page header, which are consistently present. These elements form the "shell" of the user interface for each service or feature within the portal. The header offers access to global controls,

and the configuration page for each resource provides a resource menu to facilitate navigation between features (JnHs et al., 2022).

In our project, the Azure portal was utilized as a gateway to access other Azure services. It served as a centralized platform for managing various aspects of our cloud-based resources and services. The portal's graphical interface allowed us to interact with Azure's wide array of functionalities in a more intuitive and efficient manner. By leveraging the portal's robust features, such as the customizable dashboards and the versatile menu options, we were able to streamline our workflows and manage our project's cloud resources effectively. The Azure portal's reliability and continuous availability were crucial in maintaining uninterrupted access to our cloud infrastructure, ensuring the smooth progression of our project activities.

## 4.4 Azure Resource Graph

Azure Resource Graph is a sophisticated Azure service designed to enhance Azure Resource Management (Davidsmatlak, 2023b). It offers efficient and high-performance resource exploration capabilities, enabling users to conduct complex queries at scale across various subscriptions. This functionality is vital for the effective governance of large-scale cloud environments. Azure Resource Graph allows for advanced querying of resources, including complex filtering, grouping, and sorting based on resource properties. It also facilitates iterative resource exploration based on governance requirements and enables users to assess the impact of applied policies in extensive cloud settings. Additionally, it provides the ability to track changes made to resource properties, a crucial aspect for maintaining compliance and governance in dynamic cloud environments.

One notable feature of Azure Resource Graph is its integration with the Azure portal's search bar, the browsing experience for all resources. This integration underscores its role in managing large-scale environments, offering a streamlined and cohesive experience for Azure users.

Azure Resource Graph eliminates the need for making individual calls to each resource provider, thereby simplifying access to a wide range of resource properties. This streamlined approach significantly enhances the efficiency of resource management and compliance assessment in Azure environments.

Azure Resource Graph employs the Kusto Query Language (KQL) used by Azure Data Explorer, which is pivotal for constructing queries within the service. Understanding KQL is essential for effectively utilizing Azure Resource Graph, as it enables detailed operations and functions for resource management. Users can explore resources, perform queries, and gain insights into their cloud environment's health and compliance status using this robust query language. The ability to leverage KQL within Azure Resource Graph provides a powerful means of querying and understanding the vast array of resources available in Azure (Davidsmatlak, 2023b).

In our project, Azure Resource Graph was instrumental in obtaining compliance data for our dashboard. Utilizing its capability to perform wide-reaching and detailed queries, we were able to gather comprehensive compliance data efficiently. This data was crucial for our dashboard, providing us with the necessary insights to ensure that our cloud resources were compliant with relevant policies and standards. The ease of use and the extensive capabilities of Azure Resource Graph allowed us to manage our resources effectively and maintain a high standard of governance throughout the project.

## **4.5 Azure SQL Managed Instance**

Azure SQL Managed Instance stands as a cornerstone in Microsoft's Azure cloud services, offering a fully managed database platform that significantly simplifies database administration (Neugebauer, 2023). This service provides a seamless cloud migration path for SQL Server databases, backed by the robust capabilities of a Platform as a Service (PaaS) offering. It expertly

manages critical database functions such as patching, backups, and monitoring, thereby lifting the burden of these tasks from users. Its near-perfect compatibility with the latest Enterprise Edition of SQL Server and native virtual network (VNet) implementation addresses key security concerns, making it a favorable choice for SQL Server customers transitioning to the cloud. With features like automated patching, version upgrades, and a 99.99% uptime SLA, Azure SQL Managed Instance not only enhances operational efficiency but also assures high availability and security.

Azure SQL Managed Instance is distinguished by its comprehensive feature set, supporting a wide array of SQL Server functionalities including disaster recovery, data masking, and integrated authentication. The service offers flexibility in configuring compute, memory, and storage according to specific workload demands. It is available in two service tiers – General Purpose and Business Critical – catering to different performance and high availability needs. The General Purpose tier is optimized for business applications with typical performance requirements, while the Business Critical tier is designed for applications demanding the highest levels of performance and availability, complete with local SSD storage and read-only replicas (Neugebauer, 2023).

In our project, Azure SQL Managed Instance played a pivotal role in managing our compliance data. It served as a robust and secure storage solution, ensuring that our Power BI setup could connect to it and update visualizations in real-time using DirectQuery. This integration was crucial in maintaining an up-to-date and interactive dashboard, enabling us to monitor compliance status dynamically and make informed decisions quickly. The managed instance's high availability and performance ensured uninterrupted access to compliance data, enhancing the overall efficiency and effectiveness of our project's data management strategy.

## 4.6 Citrix Virtual Desktop

Citrix Workspace, a leading-edge technology platform, has played a pivotal role in our project by redefining the concept of virtual workspaces (Wang & Citrix Staff, 2023). It integrates various applications, services, and data into a singular, cohesive user experience, which is essential in modern work environments characterized by a blend of remote and in-office work. This integration enables seamless access to resources, irrespective of the user's location or the device in use. In our project, Citrix Workspace was instrumental as a VPN for accessing State Street remotely, offering options to connect through the Virtual Desktop Infrastructure (VDI) or via our State Street laptops. This flexibility was critical in maintaining continuity and efficiency in our workflows.

Citrix Workspace stands out for its robust security measures, ensuring that sensitive data and applications are securely managed, which is a paramount concern for distributed workforces. The platform's personalization capabilities further enhance productivity by allowing users to tailor their digital workspace according to their specific needs and preferences. This adaptability is consistent across different devices and operating systems, providing a uniform experience. Another noteworthy feature of Citrix Workspace is its advanced analytics and monitoring tools, which offer valuable insights for optimizing resource utilization and identifying potential issues early on (Wang & Citrix Staff, 2023). This comprehensive approach not only simplifies the management of digital workspaces but also promotes collaboration and efficiency, thereby playing a crucial role in modernizing and streamlining organizational workflows.

## 4.7 Continuous Cloud Optimization Insights

Continuous Cloud Optimization Insights (CCO Insights) is an advanced suite of Power BI Desktop Reports, intricately designed for the monitoring and analysis of Azure Platform resources,

with added functionalities for Azure DevOps and GitHub (Suñé, 2023). Developed using Power Query M language for data extraction and transformation and Data Analysis Expressions (DAX) for advanced analytics, CCO Insights effectively harvests data from various Azure REST APIs, which serve as interfaces to Azure services, providing essential insights into Azure advisor optimizations, security alerts, and resource management strategies. Its key strength is in facilitating data-driven decision-making for cloud deployment optimization and operational efficiency enhancement. The suite features multiple dashboards, each dedicated to specific cloud platform management aspects like networking, compute, RBAC, and idle resources in the Azure Infrastructure Dashboard, and management groups, resource tagging, and policy compliance in the Azure Governance Compliance Dashboard. It also includes detailed views of project activities and contributions on GitHub and Azure DevOps, thereby aiding in better project management and team collaboration.

In our project, CCO Insights served as a foundational reference for redesigning our dashboard. Initially, we implemented it for a brief period, but we eventually ceased its usage after only a few days due to complications with its custom connectors. Despite its potential, the challenges encountered with these connectors led to our decision to abandon CCO Insights. Nevertheless, it still stands as a valuable tool for those seeking to optimize cloud infrastructure and workflows. Accompanied by a user-friendly deployment guide, CCO Insights simplifies installation and configuration processes, offering a customizable user experience. Its integration with Power BI elevates the tool's functionality, enabling users to not only visualize complex data sets but also derive actionable insights, thereby confirming CCO Insights as a vital component in cloud resource management and monitoring (Suñé, 2023).

## 4.8 Jira

Jira Software by Atlassian played a crucial role in our project, serving as a versatile project management and issue-tracking tool (Shsagir et al., 2023). Its adaptability to various management styles, including Agile methodologies like Scrum and Kanban, made it an ideal choice for our project needs. Jira Software offers customizable workflows, allowing us to tailor the processes specifically to our project requirements. Its key features that we heavily utilized include comprehensive project tracking through backlogs, sprint planning, and detailed reporting. These features, coupled with a broad range of integration capabilities with developer tools, established Jira as a centralized hub for our project management activities.

The platform's emphasis on automation was particularly beneficial in streamlining repetitive tasks, thereby enhancing efficiency and reducing the need for manual intervention. Moreover, Jira's robust scalability and security features, such as options for cloud and self-hosted deployments, SAML single sign-on, and enforced two-step verification, ensured that our project data remained secure and compliant with high standards of data security (Shsagir et al., 2023). In summary, Jira Software stood out as a versatile and powerful tool in our project, offering flexibility, robust integrations, and automation capabilities, which were instrumental in streamlining and enhancing our project management process.

## 4.9 Kusto Query Language

Kusto Query Language (KQL) is a versatile and powerful tool designed for data exploration, enabling users to discover patterns, identify anomalies and outliers, and create statistical models (Shsagir, et al. 2023). KQL uses a schema that is organized hierarchically, similar to SQL, with databases, tables, and columns. This structure facilitates a systematic approach to data analysis and

query writing. The language is accessible through the Azure Data Explorer web UI, and there are numerous tutorials available for learning how to use its common operators and functionalities. A Kusto query is essentially a read-only request that processes data and returns results. These queries are stated in plain text and follow a data-flow model that is straightforward to read, author, and automate. A typical Kusto query consists of one or more query statements, each designed to perform a specific function on the data set.

The most common type of query statement in KQL is the tabular expression statement. In this format, both the input and the output of the query are in the form of tables or tabular datasets. These statements contain various operators that start with tabular input and yield tabular output. The operators are sequenced by a pipe (|), allowing data to flow seamlessly from one operator to the next. This process is akin to a funnel where data is filtered, rearranged, or summarized at each step, leading to a refined and meaningful output at the end. An important aspect of KQL is its case sensitivity. Everything within KQL, including table names, column names, operators, and functions, is case-sensitive. This characteristic underscores the need for precision and attention to detail when writing and executing KQL queries (Shsagir, et al. 2023).

In our project, KQL played a critical role in gathering compliance data for our dashboard. By leveraging KQL's querying capabilities, we were able to efficiently process and analyze data from our Azure environment. This enabled us to extract meaningful insights and patterns related to compliance, which were crucial for the development and enhancement of our dashboard. The structured and hierarchical nature of KQL, along with its powerful data processing abilities, made it an ideal choice for handling the complex data analysis requirements of our project.



## 4.10 Microsoft Excel

Microsoft Excel, an integral component of the Microsoft Office Suite, stands as a testament to versatility and functionality in business tools (Rosenberg, 2023). Renowned for its extensive capabilities, it ranges from straightforward data entry to complex financial modeling, catering to a broad array of professional requirements. The software's user interface is designed to be intuitive and user-friendly, appealing to users at all levels, from beginners to seasoned professionals. This ease of use is further bolstered by its exceptional data analysis and visualization capabilities. Excel is particularly adept at handling large datasets, allowing users to conduct intricate calculations and generate various data visualizations effortlessly. Its proficiency in these areas renders it an indispensable tool in business contexts, where making data-driven decisions is paramount (Rosenberg, 2023).

In our project, Excel played a crucial role in manually inspecting data from Power BI. This functionality allowed us to delve deeper into the data, providing a hands-on approach to analysis and interpretation. The customization features of Excel, such as macros and Visual Basic for Applications (VBA) scripting, enable significant task tailoring and automation, thereby boosting productivity and efficiency. Its compatibility with other software, including multiple Microsoft and third-party applications, fosters smooth collaboration and data exchange. Furthermore, Excel maintains a strong emphasis on security and reliability, especially when managing sensitive information. Its continual evolution, marked by consistent updates introducing new features and enhancements, ensures its relevance and effectiveness in the ever-evolving business landscape (Rosenberg, 2023). Altogether, these characteristics firmly position Microsoft Excel as a dynamic and indispensable tool in the business arena.

## 4.11 PowerBI

Power BI, developed by Microsoft as part of its Power Platform, is a comprehensive data visualization and business intelligence tool that excels in turning complex data sets into accessible and interactive visualizations and reports (Microsoft, 2023a). Aimed primarily at data analysts and business professionals, Power BI provides an intuitive platform for creating detailed, customizable reports and dashboards. This tool is particularly adept at processing large-scale enterprise data, boasting robust capabilities in semantic modeling and advanced analytics. It efficiently handles the requirements of various users, from individual analysts to large teams, scaling as needed to support thousands of users with its enterprise-grade data processing capabilities.

One of Power BI's most notable features is its deep integration with other Microsoft services, such as Teams, PowerPoint, Excel, and the broader Power Platform. This integration significantly enhances user experience by allowing for seamless embedding and sharing of interactive reports and insights. Users can easily incorporate their Power BI findings into presentations, collaborative team platforms, and detailed spreadsheets. This cross-application functionality not only streamlines workflows but also fosters a more collaborative and efficient approach to data analysis and decision-making. The ability to embed these reports into common tools simplifies information sharing and accessibility, making it a highly effective tool for organizations that rely on Microsoft's ecosystem.

Moreover, Power BI is committed to ensuring high standards of data governance, security, and compliance, in line with Microsoft's stringent protocols. It provides end-to-end visibility over data, ensuring that organizations can manage and control their information securely. This emphasis on security is particularly important in an era where data breaches and privacy concerns are paramount. Power BI's security features include comprehensive access controls, data loss prevention,

and compliance with global standards, ensuring that sensitive data is handled with utmost care. This makes Power BI not just a powerful tool for data visualization and analysis but also a trusted platform for handling and interpreting sensitive business data. By balancing flexibility and control, Power BI positions itself as an indispensable tool in the modern data-driven landscape, where quick access to insights and secure data management is key to organizational success (Microsoft, 2023a).

In our project, Power BI emerged as a powerful tool for transforming raw data into actionable insights. As our compliance data grew in complexity, Power BI enabled us to create interactive dashboards and reports that presented compliance issues in a visually informative manner. One of the key advantages of Power BI was its ability to connect to various data sources, including KQL queries and Azure compliance data. By leveraging Power BI's data modeling and visualization capabilities, we could illustrate compliance trends, identify high-risk areas, and track progress in real-time.

## 4.12 Azure Blob Storage

Azure Blob Storage, an integral element of Microsoft's Azure cloud offerings, is tailored for handling large volumes of unstructured data, such as text and binary files (Dubey, 2023). This type of data, which does not conform to a specific model, benefits immensely from Blob Storage's optimized, scalable environment. Its key capabilities include delivering content directly to browsers, facilitating distributed file access, streaming media, logging data, and providing robust solutions for data backup, disaster recovery, and analytical processing. The adaptability of Blob Storage makes it suitable for a broad array of applications, ranging from media delivery to complex data analytics.

The architecture of Azure Blob Storage is structured around three core components: storage accounts, containers, and blobs. Storage accounts establish a unique namespace in Azure, essential

for data organization and access. Containers within these accounts serve as directories to organize blobs, which are the fundamental units of storage. Blob types vary, including block blobs for text and binary data, append blobs for logging, and page blobs for storing virtual hard drive files. Azure Blob Storage's integration with Azure Data Lake Storage Gen2 extends its functionality, combining the hierarchical file system of Data Lake Storage Gen2 with Blob Storage's features such as tiered storage, high availability, and disaster recovery capabilities. This integration is particularly beneficial for large-scale data analytics and enterprise applications (Dubey, 2023).

In the second project, the utilization of Azure Blob Storage was focused toward the end of our workflow, where it served as the storage solution for PDF compliance report files. These reports, vital for tracking compliance within the project, were securely and efficiently managed within Blob Storage, providing a reliable and scalable storage solution. The seamless access and organizational structure of Azure Blob Storage ensured that these crucial documents were readily available and well-managed, contributing significantly to the project's overall data management and compliance tracking strategy.

## **4.13 Azure Virtual Machine**

Azure Virtual Machines (VMs) are a pivotal component in Microsoft Azure's suite of scalable computing resources, designed to offer the flexibility and efficiency of cloud-based virtualization without the overhead of managing physical hardware (Nottingham, 2023). They are tailored for a range of applications, including development and testing environments, hosting cloud applications, and extending traditional data centers. Key considerations before deploying an Azure VM include selecting appropriate resource names, deciding on the geographic location for data storage, VM sizing, operating system choice, and post-launch configurations. These VMs also

provide robust availability options, such as Availability Zones for high connectivity and Virtual Machine Scale Sets for scalable, load-balanced management. The pricing model is based on VM size and operating system, with storage costs billed separately.

Azure VMs support a variety of Linux and Windows distributions, ensuring compatibility and optimization for diverse requirements. The platform's integration with cloud-init across supported Linux distributions simplifies VM deployments, particularly in scale sets. Azure Managed Disks, another integral feature, offers effortless storage management, scaling disk size and performance tiers as needed. Networking and storage solutions are comprehensive, with detailed guidance available for handling service disruptions. Azure's disaster recovery capabilities, like Azure Site Recovery and geo-redundant storage, provide additional layers of data protection and resilience (Nottingham, 2023).

In our project, Azure Virtual Machines play a crucial role by hosting the backend infrastructure. This setup allows for the automated updating of compliance reports to Azure Blob Storage, ensuring that data is consistently and securely stored. Additionally, Azure VMs support the hosting of our React front, providing a reliable and scalable platform for our application's user interface. This robust infrastructure underpins our project's ability to efficiently manage and present compliance data, demonstrating the versatility and power of Azure Virtual Machines in practical cloud computing applications.

## 4.14 Express

Express.js is a JavaScript library that significantly streamlines the handling of web routing and HTTP requests for web applications (Caceres, 2022). Building on Node.js's native HTTP library, Express.js offers an object-oriented approach to routing, making it simpler and more intuitive.

Described as a 'minimalist framework', Express.js focuses on core functionalities while leaving room for extensions through middleware, which can either be obtained from npm or created by the developer.

The primary challenge that Express.js addresses is the routing of HTTP communications in web applications. This involves managing the intricate exchange between clients (request initiators) and servers (response providers) across the internet using the TCP/IP protocol. In a typical request/response cycle, the server listens for a request, processes it, and sends back a singular, corresponding response. This cycle is fundamental to online networking, with each client request prompting a distinct server response, whether it be an HTML page, an image, or other data types.

Express.js simplifies this process by providing a streamlined method for routing requests and responses. It allows developers to easily configure routes using HTTP verbs like GET and POST, and handle the requests and responses through a callback function, which uses request and response objects. These objects provide access to all aspects of HTTP communication, such as the request body and method, or the response status and headers. Furthermore, Express.js introduces 'middleware', a powerful concept that facilitates the addition of functionalities to the app, such as logging, parsing request bodies, handling static resources, or custom error handling (Caceres, 2022).

In our project, Express.js plays a vital role in the backend server architecture. It is used for managing the server logic required for the upload of PDF files to Azure Blob Storage. Through Express.js, we efficiently handle HTTP requests for file uploads, process these requests, and execute the necessary logic to store the files in Blob Storage. This setup not only ensures smooth operation of the backend but also demonstrates the effectiveness and flexibility of Express.js in managing server-side logic in web applications.

## 4.15 React

React, a popular JavaScript library for building user interfaces, offers an intuitive approach to creating and managing UI components (Meta Open Source, 2023). These components, which can range from simple elements like buttons to complex structures like entire pages, encapsulate both the logic and appearance of the UI. In React, components are essentially JavaScript functions that return markup, commonly written in JSX, a syntax extension for JavaScript that resembles HTML. JSX is more strict than HTML, requiring tags to be closed and elements to be wrapped in a parent element for multiple JSX tags.

React's component-based architecture facilitates efficient UI development. For instance, styles are assigned using the `className` attribute, and data is displayed using JSX's curly braces to embed JavaScript expressions within the markup. React also allows for conditional rendering using standard JavaScript logic, such as `if` statements or ternary operators, and for rendering lists of data using the array's `map()` function. Each list item in React should have a unique `key` attribute for optimal rendering and updating. Event handling in React is achieved through functions within components, enabling interactive UIs (Meta Open Source, 2023).

In our project, React plays a crucial role in developing the front-end. It provides the structure for the user interface, allowing us to build a responsive and interactive application. We use React's component model to create a dynamic interface that interacts with our backend, hosting automated updates of compliance reports to Azure Blob Storage. React's efficient update and rendering system ensures our front-end stays responsive and up-to-date with the latest compliance data, contributing significantly to the overall functionality and user experience of our project.

## 4.16 Recharts

Recharts, a versatile charting library for React, offers a streamlined approach to creating visually appealing and interactive data visualizations (Recharts Group, 2023). The process begins with selecting the appropriate chart type and importing necessary components, such as `LineChart` and `Line`, to fit the specific data representation needs. The library's capability to selectively import components is beneficial for optimizing the project's size and performance. Recharts make it straightforward to generate charts using JavaScript data, usually an array of objects, thereby facilitating easy integration with various data sources.

Further customization and enhancement of the charts are achieved by adding components like `XAxis`, `YAxis`, `CartesianGrid`, and `Legend`. These additions not only enrich the visual appeal but also improve the functionality and readability of the charts. Customization extends to adjusting properties of individual components, such as margins and stroke styles, to cater to specific visualization requirements. The addition of interactive elements, such as tooltips, significantly enhances user engagement, providing detailed insights on hover and enriching the overall user experience. The ability to customize individual components, like x-axis labels, through custom render functions, showcases Recharts' flexibility and adaptability to varied visualization needs (Recharts Group, 2023).

In our project, Recharts played a pivotal role in the front-end for creating dynamic compliance history visualizations. Its seamless integration with React and extensive customization options allowed us to effectively represent complex compliance data in an intuitive and interactive manner. The library enabled us to construct charts that were not only informative but also engaging, facilitating a deeper understanding of compliance trends and patterns. The interactive features like tooltips brought an additional layer of user interaction, making the data exploration process more accessible and insightful. Recharts' comprehensive features and customization capabilities were



instrumental in enhancing the front-end user experience of our project, demonstrating its effectiveness in handling real-world data visualization challenges.

## 4.17 Vite

Vite is a modern build tool designed to enhance the development experience of web projects (You & Vite Contributors, 2023). It is composed of two major parts: a development server offering advanced features like fast Hot Module Replacement (HMR), and a build command that efficiently bundles code with Rollup for optimized production assets. Vite stands out for its opinionated yet sensible defaults, robust plugin support for frameworks and tool integration, and extensive extensibility through its Plugin API and JavaScript API with full typing support. It is tailored for modern browsers during development, assuming support for the latest JavaScript and CSS features. For production, Vite targets browsers with ES Modules, ESM dynamic import, and import.meta support, with options to accommodate legacy browsers through official plugins (You & Vite Contributors, 2023).

In the Compliance History Dashboard, Vite served as the fundamental build tool for the front-end. Its fast and efficient development server greatly streamlined the development process, providing quick feedback and updates during coding. The production build process, optimized for performance, ensured that our front-end was not only responsive but also lightweight, crucial for a smooth user experience. The ability to use plugins and extend Vite's functionality allowed us to tailor the build process to our project's specific needs, demonstrating Vite's versatility and effectiveness as a modern web development tool.

## 4.18 Communication Technologies - Discord, iOS Messaging, Microsoft Outlook, Microsoft Teams, Slack, Zoom

Throughout the project, our team utilized a variety of communication software to facilitate both internal and external communications effectively. Notably, we employed Discord for all internal team meetings, providing a versatile platform for real-time discussions and collaboration. For more informal and rapid communication within the team, as well as occasional interactions with our project sponsor, we relied on iOS Messaging. This tool enabled us to maintain a continuous and casual dialogue, essential for the swift exchange of ideas and updates.

For more formal and structured communication, particularly those involving sensitive information, Microsoft Outlook served as our primary email platform. This platform was instrumental in managing emails and scheduling meetings with various stakeholders. Meanwhile, Microsoft Teams played a dual role. It was an essential tool for sensitive communications within State Street, offering robust features that supported secure and efficient messaging crucial for handling confidential company information. Additionally, Microsoft Teams was pivotal in the version control for our PowerBI project, ensuring data integrity by providing a centralized and secure environment for file sharing and collaboration.

Slack played a pivotal role in our communication with advisors, particularly within the “Fall23 Fintech MQPs” server. This platform, in conjunction with Microsoft Outlook, facilitated a seamless exchange of information and feedback with our academic advisors, ensuring a cohesive advisory process. Lastly, Zoom was the preferred choice for video meetings with both our advisors and State Street personnel. The platform's reliable video conferencing capabilities allowed us to conduct productive and engaging virtual meetings, vital for maintaining strong collaborative relationships throughout the project.

## 5. Software Requirements

### 5.1 Software Requirements Gathering Strategy

In the early stages of project planning, our approach to gathering software requirements was largely shaped through informal discussions with our project sponsor. His unique position of having sole control over the project meant that our daily meetings with him were not just routine check-ins but crucial strategy sessions. These interactions played a pivotal role in refining the project's scope and direction, especially considering the project's dynamic nature.

Initially, the scope and technological stack were fluid, evolving significantly through the first four weeks. We started implementation of the first project in the third week, due to delays in setup and configuration. We navigated through a landscape of changing requirements and shifting priorities. This period was marked by a high degree of adaptability from our team, as we continuously integrated feedback from our project sponsor into our development process. The discussions during these meetings were instrumental in narrowing down our focus to what was practically achievable, balancing the educational and business value.

Throughout this phase, our team faced challenges typical of rapidly evolving projects, including aligning our development pace with the changing requirements and ensuring that our work remained aligned with our project sponsor's vision. This necessitated an agile and responsive approach, allowing us to make adjustments on the fly and address unforeseen issue blockers effectively. Our engagement with our sponsor was not just about gathering requirements; it was a collaborative effort in steering the project towards a product that would meet the organization's needs while providing us with a rich learning experience. By the end of this initial phase, we had had

a clear, albeit evolving, path for our project, thanks to the constant guidance and feedback from our sponsor.

As we finalized the first project, our team embarked on the development of the Compliance History Dashboard, a separate but related endeavor. This second project began concurrently with the initial dashboard but with a distinct focus and developmental pathway. Operating outside the conventional State Street environment, this project allowed us the freedom to experiment with a different technological stack and Azure services not available in their standard cloud setup. The early stages were marked by the creation of a similar data model to that of the Governance Compliance Dashboard, yet with a key difference: the reliance on mock data for testing and development. This approach was crucial in shaping the dashboard's functionality, as it enabled us to simulate real-world scenarios within a controlled, experimental framework.

## **5.2 Functional and Non-functional Requirements**

In the realm of functional requirements for both projects, paramount importance is placed on data refresh, accurate data, and the ability to perform meaningful visualizations and metrics. Data refresh is critical, especially considering the dynamic nature of compliance issues; old data quickly becomes obsolete, necessitating a system that ensures continuous updates and relevance. Accuracy of data is another cornerstone, as it directly influences the integrity and reliability of the entire system. This requirement is crucial for making informed decisions and maintaining the credibility of the project. Additionally, the capacity to conduct meaningful visualizations and metrics is integral. This feature transforms raw data into actionable insights, enabling stakeholders to discern patterns, trends, and anomalies, thereby facilitating strategic planning and operational efficiency.

On the non-functional front, data security stands as a critical barrier, ensuring that sensitive information remains protected from unauthorized access. This aspect of security is especially

pertinent in the financial sector where data breaches can have catastrophic consequences.

Performance is another key non-functional requirement, where the system must operate efficiently under varying loads, maintaining responsiveness and reliability. Scalability, closely tied to performance, is essential for the system to adapt to growing data volumes and user demands without compromising on performance or user experience. Cost considerations also play a significant role in non-functional requirements, impacting choices related to technology, resources, and long-term maintenance of the project. Customization and personalization are vital for tailoring the system to meet specific user needs and preferences, enhancing user engagement and satisfaction. Accessibility is another crucial aspect, ensuring that the system is usable by people with a wide range of abilities and disabilities, thus broadening the user base and adhering to inclusivity principles. Lastly, user-friendliness is important; a system, regardless of its technical sophistication, must be intuitive and easy to use to ensure widespread adoption and effective utilization. These requirements are not just quality indicators but are fundamental to the long-term viability of the project.

### 5.3 User Stories and Epics

Sprint	User Story	Points
<b>Epic: Setup and Configuration</b>		
0	Develop project scope	N/A
0	Familiarize with Graph API	N/A
0	Familiarize with open source CCO Insights	N/A
0	Familiarize with Power BI	N/A

Sprint	User Story	Points
0	Familiarize with structure of existing dashboard	N/A
0	Data gathering with Custom Connector	N/A
<b>Epic: Compliance Data Visualization</b>		
<b>Section: Filters</b>		
1	Add Management Group Checkbox Filter	3
1	Add Subscription Checkbox Filter	1
1	Add Policy Assignment Name Checkbox Filter	1
1	Add Policies Filter by Policies	1
1	Add Environment Dropdown Filter	2
1	Add Policy Effect Dropdown Filter	2
1	Add Severity Dropdown Filter	2
1	Add Compliance State Dropdown Filter	1
1	Add APP_CODE Dropdown Filter	2
<b>Section: Primary Visualizations</b>		
1	Add Policy Compliance Pie Chart	8

Sprint	User Story	Points
1	Add Resource Compliance Pie Chart	8
1	Add Top 10 Non-Compliant Policies Clustered Bar Chart	5
1	Add Top 10 Non-Compliant RGs Clustered Bar Chart	5
1	Add List of Policies Table	5
<b>Section: Secondary Visualizations</b>		
1	Add Resource Compliance Percentage	3
1	Add Policy Compliance Percentage	5
1	Risk Score Visualization	5
1	Add Total Resource Count	1
1	Add Total Policy Count	3
1	Add Total Assignment Count	5
1	Add Total Initiative Count	5
<b>Section: Miscellaneous</b>		
1	Move List of Policies Table to new page	2
1	Maintain filters between the two pages	3

Sprint	User Story	Points
3	Plan for Connecting Power BI Dashboard to Data Source	8
<b>Epic: Compliance Data Automation</b>		
2	Get access to Azure Automation Account	2
2	Get access to SQL MI	2
2	Create a new database in the SQL MI	1
2	Add sqlServer Module to Automation Account	3
3	Add the data into the SQL MI	8
2	Run Resource Graph Query in the runbook	5
3	Test the runbook	2
2	Access the SQL MI from PowerShell	8
2	Schedule the runbook weekly	2
2	Add Azure Authentication to runbook	5
<b>Epic: Compliance History Dashboard</b>		
2	Create a project plan	6
2	Configure Git repository	1



Sprint	User Story	Points
2	Configure React and Node.js environment	1
2	Generate Dummy Compliance Data	4
2	Visualize Compliance History Chart	6
3	Style the charts and website	8
3	Create report template	8
3	Style the report	6
3	Create a watermark on the reports	4
3	Add export report in PDF format	8
3	Create Azure Blob Storage for automated reports	4
3	Auto generate compliance reports every week	8
3	Refactor code	6
3	Deploy application into Azure VM	8

Table 5.3: User Stories and Epics

## 6. Design

### 6.1 Governance Compliance Dashboard




Figure 6.1: Project 1 High-Level Project Architecture

The Governance Compliance Dashboard project involves a well-orchestrated architecture utilizing Microsoft technologies. At the core of the data retrieval strategy is the use of Microsoft Graph API, a unified interface allowing programmable access to a diverse range of data from Microsoft 365 services. This API serves as the primary source for the data that populates the dashboard.

To ensure continuous updating of the data in defined intervals, Azure Automation is utilized. A runbook inside Azure Automation is created to run a PowerShell script to autonomously run at defined intervals to execute queries against Microsoft Graph API to retrieve data from the Azure Resource Groups and Policy from Azure Portal. The retrieved data is then placed inside of the SQL Managed Instance. This periodic execution not only facilitates automated updates but also allows for the parsing and transformation of the received data. This step is crucial for shaping the raw data into a structured format suitable for storage and subsequent analysis.

Azure SQL Managed Instance (MI) is leveraged as the central repository for the data. This fully managed SQL Server instance in Azure offers high availability and scalability, providing a robust solution for storing the parsed and transformed data from the Microsoft Graph API. The choice of Azure SQL MI reflects a commitment to a scalable and reliable storage solution for governance compliance data.

Power BI is seamlessly integrated into this architecture, using DirectQuery to establish a connection with Azure SQL MI. DirectQuery enables Power BI to send queries directly to the underlying data source, ensuring that reports and dashboards reflect the most up-to-date information stored in Azure SQL MI. This approach is particularly suitable for scenarios with large datasets where importing data into the Power BI dataset might not be optimal. Within the Power BI environment, Power Query is employed for further data preprocessing. This powerful tool allows for the creation of new fields, calculations, and additional transformations. Power Query acts as a critical component in refining the data before it is used for visualization and analysis.

In the final stage, visualizations, calculations, and analyses are created in Power BI based on the preprocessed and transformed data. Users interact with the dashboard and the data is dynamically fetched from Azure SQL MI using DirectQuery, ensuring that the dashboard provides an accurate view of governance compliance metrics.

## 6.2 Compliance History Dashboard

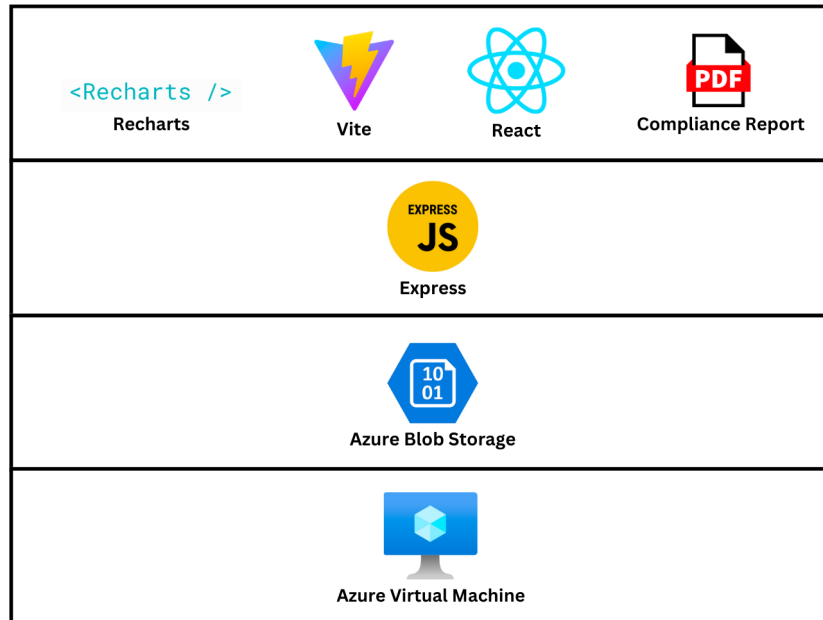


Figure 6.2: Project 2 High-Level Project Architecture

The Compliance History Dashboard project serves as an extension to the Governance Compliance Dashboard project which adopts a meticulously orchestrated architecture, employing a combination of technologies within the Microsoft ecosystem. This project was developed outside of the State Street environment to explore the feasibility and potential enhancements to the Governance Compliance Dashboard using different technologies and Azure services that are not enabled on State Street’s cloud environment. A similar data model was used in this project as in the Governance Compliance Dashboard project; however, mock data was generated due to being outside the State Street environment. The dashboard was created on a full-stack website.

The front-end of the website was developed using React, a JavaScript library for building user interfaces, and Recharts, a visualization library. React serves as the foundational framework for the front-end, offering a modular and efficient approach to organizing user interface components.

Vite was also used to enhance the development experience by supporting the latest ECMAScript features and providing quick development server start-up. Recharts, on the other hand, facilitates the creation of dynamic and visually engaging compliance history visualizations, including bar charts and line charts, showcasing the historical data of compliant and noncompliant policies.

Express, a minimal and flexible Node.js web application framework, was used for the backend server logic, and constitutes the backend server responsible for uploading PDF compliance reports to Azure Blob Storage.

Azure Blob Storage was employed as the storage solution for the generated PDF compliance report files, offering a scalable and secure environment. The automation of compliance report updates is executed through an Azure Virtual Machine, which hosts the web application. This virtual machine ensures the periodic and automated updates of compliance reports to Azure Blob Storage, enhancing the reliability and scalability of the system.

### 6.3 Data Model Entity Relationship Diagram

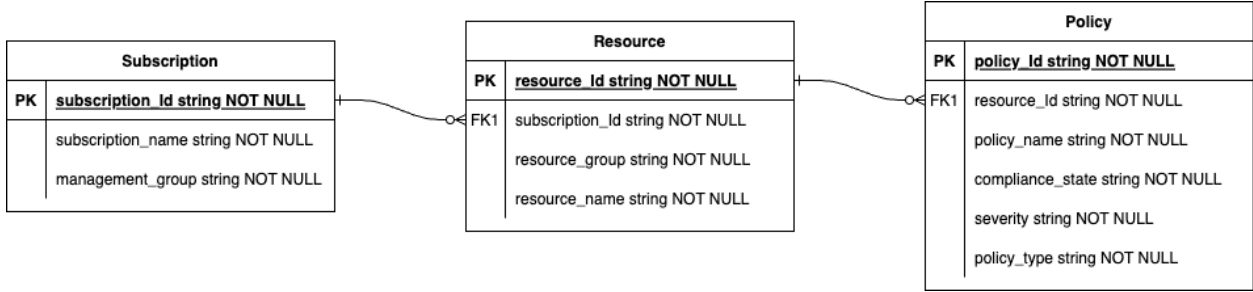


Figure 6.3: Entity Relationship Diagram

The Entity-Relationship Diagram (ERD) for both projects illustrates the relationships and associations between three key tables used to create our data model: Resources, Policies, and

Subscriptions. This represents a high-level overview of the ERD we utilized in the project, where certain columns have been omitted to ensure the exclusion of any confidential information.

### 6.3.1 Resources Table

The Resources table serves as a central repository for information about various Azure resources within the State Street environment. The table includes the following attributes which are all strings:

- Primary Key (PK): resource\_id
- Foreign Key (FK): subscription\_id (linked to the Subscriptions table)
- resource\_group: Denotes the resource group to which the resource belongs.
- resource\_name: Represents the name of the resource.

The relationship between the Resources table and the Subscriptions table is established through the foreign key subscription\_id. This signifies that a subscription can have zero or many associated resources, but each resource is associated with exactly one subscription.

### 6.3.2 Policies Table

The Policies table captures information related to compliance policies applied to specific Azure resources. This table includes the following attributes which are all strings:

- Primary Key (PK): policy\_id
- Foreign Key (FK): resource\_id (linked to the Resources table)
- policy\_name: Describes the name of the policy.
- compliance\_state: Represents the compliance status of the resource concerning the policy (compliant or non-compliant).

- severity: Indicates the severity level associated with the policy violation (low, medium, high, or critical).
- policy\_type: Specifies the type or category of the compliance policy (built-in, custom, or other).

The relationship between the Policies table and the Resources table is established through the foreign key resource\_id. This indicates that a resource can have zero or many associated policies, but each policy is associated with exactly one resource.

### 6.3.3. Subscriptions Table

The Subscriptions table holds information about Azure subscriptions within the State Street environment. This table includes the following attributes which are all strings:

- Primary Key (PK): subscription\_id
- subscription\_name: Represents the name or identifier of the subscription.
- management\_group: Indicates the management group to which the subscription belongs.

There is a one-to-zero-or-many relationship between the Subscriptions table and the Resources table. This means that a subscription can have zero or many associated resources, but each resource is associated with exactly one subscription.

The ERD visually represents the hierarchical and relational structure of the three tables, illustrating how Azure resources are associated with subscriptions and how compliance policies are linked to specific resources. This design facilitates the organization and retrieval of data for the development of both projects.

## 7. Software Development

Our project embraced the agile scrum methodology, a decision that profoundly influenced our workflow and project management strategies. Agile scrum, known for its flexibility and iterative nature, was not just a methodology but a guiding philosophy throughout our project. We adopted one-week sprints, each spanning from Monday to Friday. This structured yet adaptable approach allowed us to frequently reassess and realign our objectives, ensuring that our project adapted to emerging challenges and evolving requirements. The weekly rhythm fostered by these sprints proved invaluable, instilling a sense of discipline while accommodating the dynamic nature of our work.

At the start of each week, our sprint planning sessions were instrumental in charting the course for the days ahead. These Monday meetings transcended typical administrative gatherings; they were strategic sessions where we collectively defined our weekly goals, delegated tasks, and established clear, measurable targets. This meticulous planning was crucial in creating a roadmap for each sprint, clarifying roles, and setting realistic expectations. It ensured that every team member was not only aware of their responsibilities but also aligned with the team's collective objectives.

Following the sprint planning, daily sponsor meetings, initially set at 9:30 AM and later moved to 10:00 AM, became a cornerstone of our daily routine. In these meetings, our interaction with the primary contact provided essential insights that aligned our project's direction with the sponsor's vision. His feedback and guidance were vital in steering our project on the right path and ensuring that we met our sponsor's expectations. Immediately following the sponsor meetings we had daily standups conducted on Discord, which were integral to our internal communication. These standups went beyond routine check-ins; they were vibrant, interactive sessions where team members shared their previous day's achievements, outlined their plans for the day, and discussed any impediments in their path. This daily practice nurtured a culture of openness and collective



problem-solving, encouraging team members to support each other and tackle challenges collaboratively.

In addition to our verbal communications, we maintained a rigorous discipline of compiling detailed daily reports. These documents were far more than mere records; they were comprehensive accounts that captured the essence of each day's progress. They detailed the outcomes of meetings, new information acquired, accomplishments of the day, and future meeting agendas. These reports served as a chronological narrative of our project, providing a valuable resource for tracking progress, reflecting on our journey, and guiding future decisions.

Our engagement with external advisors, particularly during the weekly check-ins every Tuesday at 3:30 PM, offered an invaluable external perspective. These sessions were critical for obtaining expert advice, addressing concerns, and ensuring our project remained on a trajectory toward success. The importance of face-to-face interaction was recognized through our Thursday in-person collaborations at the Quincy State Street Office. These sessions provided a platform for more intensive teamwork and problem-solving, leveraging the benefits of direct communication and collaboration.

The week's activities culminated in sprint retrospectives every Friday evening. These retrospectives were not just procedural closures; they were reflective sessions where we collectively evaluated the week's achievements, identified areas for improvement, and strategized for future sprints. This continuous cycle of planning, execution, and reflection, guided by the principles of agile scrum, was fundamental in driving our project forward. It ensured efficient use of time, effective communication, and consistent alignment with our strategic goals, ultimately contributing to the project's successful outcome.

Monday	Tuesday	Wednesday	Thursday	Friday
Sponsor Check-In	Sponsor Check-In	Sponsor Check-In	In the office	Sponsor Check-In
Standup	Standup	Standup		Standup
	Advisor Check-In			Sprint retrospective

Figure 7.1: Weekly Overview

## 7.1 Pre-sprint Setup and Configuration

In the initial stages of our projects, the team encountered a series of challenges that hindered the full application of Agile methodologies. This period, often referred to as pre-sprint onboarding and configuration, was crucial for laying the groundwork for future progress. Despite these hurdles, significant strides were made in aligning with the compliance standards of State Street.

The pre-sprint phase was characterized by a focus on administrative and preparatory tasks rather than software development. Key activities during this phase included completing essential paperwork like fingerprinting, background checks, and I-9 forms. These steps were vital to ensure adherence to State Street’s compliance standards. In addition, the team engaged in numerous compliance training sessions, which were instrumental in acquainting members with the organization's standards and expectations. This phase was not only about meeting regulatory requirements but also about integrating the team into the company's ecosystem. Acquiring State Street ID badges, laptops, and access to critical resources like the Virtual Desktop Infrastructure (VDI) and JIRA platform was a significant part of this integration process.

A notable highlight of this phase was our journeys into the Quincy State Street office. We met our primary contact for the first time in person, as well as other members of the policy compliance team. It provided an opportunity for face-to-face interactions, which was invaluable in establishing stronger connections and enhancing collaboration.

During this period, the team also began to incrementally integrate the Governance Compliance Dashboard's goals and objectives. While our initial focus was predominantly on onboarding and compliance, we progressively started to delve into the core aspects of our project. Regular meetings and interactions with our project sponsor and the policy team allowed us to understand the intricacies of this dashboard which we were tasked with. Our understanding of the project's scope and requirements began to solidify, setting the stage for more focused and structured Agile implementation in subsequent phases.

Despite the absence of software development tasks in the pre-sprint period, the team focused on advancing other aspects of the project, particularly those related to administrative tasks and defining the project's scope. The overarching goal was the swift removal of blockers to facilitate a smooth transition into the project. This proactive approach was crucial in maintaining the momentum of the project.

## 7.2 Sprint One

### 7.2.1 Sprint Backlog

Status	Story Owner	User Story	Points
<b>Epic: Compliance Data Visualization</b>			
COMPLETE	Ryan Kornitsky, Will Huang	Add Management Group Checkbox Filter	3
COMPLETE	Ryan Kornitsky, Will Huang	Add Subscription Checkbox Filter	1
CARRIED OVER	Ryan Kornitsky, Will Huang	Add Policy Assignment Name Checkbox Filter	1
COMPLETE	Ryan Kornitsky, Will Huang	Add Policies Filter by Policies	1
COMPLETE	Ryan Kornitsky, Will Huang	Add Environment Dropdown Filter	2
COMPLETE	Ryan Kornitsky, Will Huang	Add Policy Effect Dropdown Filter	2
COMPLETE	Ryan Kornitsky, Will Huang	Add Severity Dropdown Filter	2

Status	Story Owner	User Story	Points
COMPLETE	Ryan Kornitsky, Will Huang	Add Compliance State Dropdown Filter	1
COMPLETE	Ryan Kornitsky, Will Huang	Add APP_CODE Dropdown Filter	2
CARRIED OVER	Ryan Kornitsky, Will Huang	Add Resource Compliance Percentage	3
CARRIED OVER	Ryan Kornitsky, Will Huang	Add Policy Compliance Percentage	5
CARRIED OVER	Ryan Kornitsky, Will Huang	Risk Score Visualization	5
COMPLETE	Ryan Kornitsky, Will Huang	Add Total Resource Count	1
COMPLETE	Ryan Kornitsky, Will Huang	Add Total Policy Count	3
CARRIED OVER	Ryan Kornitsky, Will Huang	Add Total Assignment Count	5
CARRIED OVER	Ryan Kornitsky, Will Huang	Add Total Initiative Count	5

Status	Story Owner	User Story	Points
COMPLETE	Ryan Kornitsky, Will Huang	Add Policy Compliance Pie Chart	8
COMPLETE	Ryan Kornitsky, Will Huang	Add Resource Compliance Pie Chart	8
COMPLETE	Ryan Kornitsky, Will Huang	Top Non-Compliant Policies Clustered Bar Chart	8
COMPLETE	Ryan Kornitsky, Will Huang	Add List of Policies Table	5
COMPLETE	Ryan Kornitsky, Will Huang	Move List of Policies Table to new page	2
COMPLETE	Ryan Kornitsky, Will Huang	Maintain filters between the two pages	3
<b>Epic: Paper</b>			
COMPLETE	Ryan Saklad	Abstract	1
COMPLETE	Ryan Saklad	Executive Summary	2
COMPLETE	Ryan Saklad	Introduction (Section 1)	3
CARRIED OVER	Ryan Saklad, Will	Background (Section 2)	6

Status	Story Owner	User Story	Points
	Huang		
COMPLETE	Ryan Saklad	Methodology (Section 3)	2
CARRIED OVER	Ryan Saklad	Software Development Environment (Section 4)	10
CARRIED OVER	Ryan Saklad	Software Requirements (Section 5)	3
COMPLETE	Ryan Saklad, Ryan Kornitsky	Design (Section 6)	6
CARRIED OVER	Ryan Saklad	Software Development (Section 7)	7
<b>Total Points Completed</b> (Maximum: 116)			66

Table 7.2: Sprint 1 Backlog

### 7.2.2 Sprint Summary

Reflecting on our latest sprint, we have made substantial strides, particularly with the Governance Compliance Dashboard, which is now around 95% complete. This achievement was largely due to finally obtaining access to the necessary resources we had requested earlier. Concurrently, we have also advanced significantly in developing the paper. A notable shift in our approach this sprint was the incorporation of agile methodologies, utilizing the Jira board for sizing stories and enhancing project management. However, we did encounter a significant challenge: a number of Azure services crucial for the data automation aspect of the Governance Compliance

Dashboard are not enabled in our current department. This limitation has posed a hurdle that we need to navigate around, potentially requiring us to rethink certain aspects of our project or seek alternative solutions. On the front of team dynamics and logistics, an area identified for improvement is the scheduling and punctuality of our internal meetings. While this was the only notable concern in this regard, addressing it will be crucial for maintaining the momentum and efficiency we have built up. Overall, this sprint has been a blend of significant achievements and valuable learning experiences. The progress made on key components like the dashboard and the paper, coupled with the adoption of agile practices, marks a positive trajectory for our project. The challenges and areas for improvement identified provide us with clear targets for enhancement in our next sprint.

### 7.3 Sprint Two

#### 7.3.1 Sprint Backlog

Status	Story Owner	User Story	Points
<b>Epic: Compliance Data Visualization</b>			
COMPLETE	Ryan Kornitsky, Will Huang	Add Policy Assignment Name Checkbox Filter	1
COMPLETE	Ryan Kornitsky, Will Huang	Add Resource Compliance Percentage	3
COMPLETE	Ryan Kornitsky,	Add Policy Compliance Percentage	5



Status	Story Owner	User Story	Points
	Will Huang		
COMPLETE	Will Huang	Risk Score Visualization	5
COMPLETE	Ryan Kornitsky, Will Huang	Add Total Assignment Count	5
COMPLETE	Ryan Kornitsky, Will Huang	Add Total Initiative Count	5
<b>Epic: Compliance Data Automation</b>			
COMPLETE	Will Huang	Get access to Azure Automation Account	2
COMPLETE	Will Huang	Get access to SQL MI	2
COMPLETE	Will Huang	Create a new database in the SQL MI	1
COMPLETE	Will Huang	Add sql Server Module to Automation Account	3
CARRIED OVER	Will Huang	Add the data into the SQL MI	10
CARRIED OVER	Will Huang	Access the SQL MI from PowerShell	8
COMPLETE	Will Huang	Run Resource Graph Query in the runbook	5
CARRIED OVER	Will Huang	Test the runbook	2

Status	Story Owner	User Story	Points
COMPLETE	Will Huang	Add Azure Authentication to runbook	5
<b>Epic: Paper</b>			
COMPLETE	Ryan Saklad, Will Huang	Background (Section 2)	6
COMPLETE	Ryan Saklad	Software Development Environment (Section 4)	10
COMPLETE	Ryan Saklad	Software Requirements (Section 5)	3
COMPLETE	Ryan Saklad	Software Development (Section 7)	7
COMPLETE	Ryan Kornitsky	Assessment (Section 8)	4
COMPLETE	Ryan Saklad, Ryan Kornitsky	Business and Risk Management (Section 9)	10
COMPLETE	Ryan Kornitsky	Future Work (Section 10)	3
COMPLETE	Ryan Saklad	Conclusion (Section 11)	1
COMPLETE	Will Huang	References	4
COMPLETE	Ryan Saklad, Will Huang	Appendix A	1
COMPLETE	Ryan Kornitsky	Acknowledgements	1

Status	Story Owner	User Story	Points
<b>Epic: New Project</b>			
COMPLETE	Ryan Kornitsky	Create a project plan	6
COMPLETE	Ryan Kornitsky	Configure Git Repository	1
COMPLETE	Ryan Kornitsky	Configure React and Node.js Environment	1
COMPLETE	Ryan Kornitsky	Generate Dummy Compliance Data	2
COMPLETE	Ryan Kornitsky	Visualize Compliance History Chart	6
<b>Total Points Completed</b> (Maximum: 128)			108

Table 7.3: Sprint 2 Backlog

### 7.3.2 Sprint Summary

During sprint two, our team's adaptability and multitasking prowess were prominently displayed as we juggled advancing our new project alongside our ongoing ones. A considerable portion of our time had been devoted to the paper, transforming it from a state of significant delay to near completion, pending only minor revisions. This marked a crucial shift in our focus, demonstrating our capacity to recover and progress rapidly on critical tasks. The initiation of the second project had added a complex layer to our workload, further testing our ability to manage multiple priorities efficiently. However, this also presented challenges, particularly in balancing resources and attention between the two projects. The interplay between these tasks required careful

navigation to ensure neither was neglected. In terms of team dynamics, punctuality and strengthened communication with WPI advisors had emerged as key areas needing improvement. Enhancing these aspects was vital for maintaining our team's efficiency and cohesion. Overall, sprint two was a journey of significant growth, learning, and adaptation, with a clear roadmap for sprint three.

## 7.4 Sprint Three

### 7.4.1 Sprint Backlog

Status	Story Owner	User Story	Points
<b>Epic: Compliance Data Visualization</b>			
COMPLETE	Will Huang	Plan for Connecting Power BI Dashboard to Data Source	8
<b>Epic: Compliance Data Automation</b>			
DEPRECATED	Will Huang	Add the data into the SQL MI	8
DEPRECATED	Will Huang	Access the SQL MI from PowerShell	8
DEPRECATED	Will Huang	Test the runbook	2
DEPRECATED	Will Huang	Schedule the runbook weekly	2
<b>Epic: Paper</b>			
COMPLETE	Ryan Saklad	Update rough draft for comments	4

Status	Story Owner	User Story	Points
COMPLETE	Ryan Saklad, Ryan Kornitsky, Will Huang	Update Section 7 for end of sprint 2 and sprint 3	1
COMPLETE	Ryan Saklad	Update Abstract, Section 1 for second project	3
COMPLETE	Ryan Saklad	Update Section 4 for second project	3
COMPLETE	Ryan Kornitsky	Update Section 6 for second project	3
COMPLETE	Ryan Saklad, Ryan Kornitsky	Update Section 8 for second project	2
COMPLETE	Ryan Saklad	Update Section 9 for second project	4
COMPLETE	Ryan Saklad	Update Section 10 for second project	2
COMPLETE	Ryan Saklad	Update Section 11 for second project	1
COMPLETE	Ryan Saklad	Update Appendix A for second project	1
<b>Epic: Compliance History Dashboard</b>			
COMPLETE	Ryan Kornitsky	Style the charts and website	8
COMPLETE	Ryan Kornitsky	Create report template	8
COMPLETE	Ryan Kornitsky	Style the report	6

Status	Story Owner	User Story	Points
COMPLETE	Ryan Kornitsky	Create a watermark on the reports	4
COMPLETE	Ryan Kornitsky	Add export report in PDF format	8
COMPLETE	Ryan Kornitsky	Create Azure Blob Storage for automated reports	4
COMPLETE	Ryan Kornitsky	Auto generate compliance reports every week	8
COMPLETE	Ryan Kornitsky, Will Huang	Refactor code	6
COMPLETE	Ryan Kornitsky, Will Huang	Deploy application into Azure VM	8
<b>Total Points Completed</b> (Maximum: 97)			75

Table 7.4: Sprint 3 Backlog

### 7.4.2 Sprint Summary

Sprint 3 was a vital period for our team, marking the successful completion of our deliverables, including both project 1 and project 2, as well as the finalization of our paper. The team's diligence and coordinated efforts culminated in products that we are collectively satisfied with, demonstrating our capability to bring complex projects to fruition. During this sprint, we focused on building a robust foundation for the data automation aspect of project 1. However, a new version of Power BI was recently released that has built-in Azure Resource Graph compatibility

with Power BI. We explored getting access to this new version of Power BI, however, State Street's system prevented us from getting the latest version. This development, while extremely beneficial, also meant that much of the work we had done earlier in the sprint became obsolete. It highlighted a key aspect of working with evolving technologies: the importance of staying adaptable and receptive to new tools and methodologies that can enhance our work, even if it means revisiting or discarding previous efforts. Reflecting on our journey, this sprint emphasized the importance of agility in adapting to new tools and methods, as well as the critical role of effective communication in managing project timelines. The completion of our primary goals is a testament to our team's hard work and dedication, while the challenges faced provide valuable lessons for future endeavors. As we move forward, we aim to integrate these learnings to enhance our efficiency and effectiveness in subsequent projects.

## 8. Assessment

### 8.1 Business Learnings

Engaging in a collaboration with State Street offered our team a deep dive into the intricacies of corporate finance, providing valuable insights into industry culture, operations, and team project management. The experience allowed us to apply knowledge gained throughout the term to develop a meaningful project and establish connections within the company.

Despite significant accomplishments, we initially struggled to develop a firm project plan with State Street early during the PQP term. Project scope issues continued to arise throughout the first half of the MQP term, especially as we found ourselves navigating without clear direction, hindering our ability to complete more technical project work than anticipated. Communication within a large financial company often proved challenging, necessitating a closer relationship with sponsors and leading us to refine our Minimum Viable Product (MVP) through more inquisitive interactions.

The introduction of the second dashboard project over halfway through the MQP term added a new dimension to our work. This proof of concept, an independent React application for detailed visualizations of historical compliance data, represented an exploratory step in the evolving landscape of financial services data management and governance. Though not utilizing actual State Street data, this development underscored the potential benefits and technical requirements of such a tool, providing a foundation for future projects with real data and resource integration.

With these challenges in mind, it was vital that when our MQP began, we needed to acknowledge when collective help was required. Setting up a regular meeting schedule to foresee technical challenges was essential, avoiding the need to address problems only when they occurred.



State Street's culture of collaboration, where employees readily shared technical expertise, played a pivotal role in overcoming obstacles. For instance, daily meetings with our project sponsor were invaluable in pushing us ahead of any blockers that we faced where he would generously assist us with any problems or reach out to others to find the help we needed. Furthermore, we often were in collaboration with State Street's international team members, particularly those in India, who make up a large portion of the compliance team, which contributed significantly to our project's success. Challenges emerged due to time zone differences, requiring adjustments for overlapping availability.

During this project, we largely worked in a remote setting. While remote work conditions were manageable, we felt it was imperative for both our own learning experiences to interact with the compliance team in person and for the development of the projects to commute into the office in Quincy once a week. Although State Street's remote communication has been excellent, without this in-person interaction, the synergy and collaborative spirit essential for tackling the challenges faced in developing the dashboard may not be as effectively fostered. On the other hand, commuting from Worcester to Quincy more than once a week would have resulted in a lot of lost time working on the paper and the projects due to commuting time, so we felt commuting to the office once a week was perfect.

In the end, both projects marked significant strides, delivering an invaluable tool to the compliance team to monitor their cloud compliance issues in one place and demonstrating what could be developed in the future. This certainly imparted great technical learnings and achievements to us. In hindsight, going in the MQP with a goal of what we want to get out of it would have aided tremendously in pushing along the project planning phase and likely would have eliminated a lot of blockers. Having a high-level weekly plan of our project life cycle before development would have also benefited us greatly in minimizing blockers and getting access since we would roughly know

which resources or technologies we would need to use far in advance. The cooperative work environment at State Street facilitated efficient learning and problem-solving, underlining the importance of collaboration as industries evolve with new technologies.

## 8.2 Technical Learnings

### 8.2.1 Azure Services

In order to explore possible project ideas to work on for project 1 and for additional ways to expand the project in the future, we first needed to understand a handful of different services that Azure offers. It was critical to understand how those services interacted with other services. Throughout this project, we researched ten different Azure services to use for prospective projects and project extensions but ultimately used the following Azure services:

- Azure Automation
- Azure SQL Managed Instance
- Azure Blob Storage
- Azure Virtual Machines

Mastering these Azure resources equip us with the skills to effectively utilize Azure tools and tailor them to address future challenges. We now grasp the structural composition of various cloud environments, enabling us to easily comprehend the behavior of upcoming environments.

### 8.2.2 Microsoft Graph API and Azure Resource Graph Explorer

The Governance Compliance Dashboard demanded a deep understanding of the Microsoft Graph API and Azure Resource Graph Explorer to extract data from policy, resources, and subscription tables. An understanding of databases and ERD diagrams was also extremely helpful

when extracting this data, for we better understood the cardinality between these three tables and how to join them together. This technical choice was vital to ensuring the accuracy of compliance metrics presented in the new dashboard. Navigating and comprehending the intricacies of these tools became a foundational aspect of the project, requiring exploration into effective data retrieval strategies. The challenge lay in adapting to updates and changes in Azure services to maintain the dashboard's relevance, but success was achieved through continuous learning and staying informed about evolving technologies. Additionally, diving through a huge number of possible tables in Azure Resource Graph Explorer to retrieve information to use in our data model proved to be a challenge. Success came through a combination of continuous learning and collaboration with various members of the compliance team.

### **8.2.3 Kusto Query Language (KQL)**

The mastery of KQL was fundamental to constructing effective queries for Azure Resource Graph Explorer in the development of the Governance Compliance Dashboard. KQL's syntax and logic became focal points of the team's learning journey for data retrieval and the creation of the data model. Challenges arose in crafting queries that optimized data retrieval while ensuring precision. Joining together numerous tables and retrieving a lot of information from each table made us heavily consider the readability of our query and adjust accordingly. Success was achieved through hands-on experimentation, collaborative discussions, and referencing KQL documentation.

### **8.2.4 Entity Relationship Diagram (ERD)**

Understanding and creating an Entity Relationship Diagram (ERD) was a crucial aspect of the projects for data modeling. The ERD visually represented the relationships between different entities in the data model, providing a clear overview of how various tables and fields are

interrelated. Constructing an ERD prior to writing the KQL query enabled us to consider the cardinality of the tables, how to join all the tables, and which columns from each to include. This resulted in minimal time writing the query and retrieving the minimal amount of information needed, greatly improving readability. Challenges included identifying the appropriate entities and defining relationships accurately. Success was achieved by collaborating with members of the compliance team to understand the data structure, finding all the necessary data that was needed for the dashboard visualization, and designing an ERD that served as a foundational blueprint for building a robust data model in Power BI.

### **8.2.5 Power BI and Power Query for Data Preprocessing**

Leveraging Power BI and Power Query for the Governance Compliance Dashboard was crucial for data preprocessing and creating meaningful visualizations. The decision to create new fields using Power Query allowed for the development of complex calculations and the transformation of raw data into insightful metrics, such as the risk score and counts of compliant resources and policies. The challenge here was to strike a balance between achieving accuracy in calculations and maintaining a streamlined data preprocessing workflow. Success was achieved by delving into the capabilities of Power BI and Power Query, empowering the team to translate complex compliance metrics into visually appealing insights.

### **8.2.6 Power Query M Formula Language**

Power Query M formula language was indispensable for the development of the Governance Compliance Dashboard for efficient data processing and transformation within Power BI. It shaped and refined the data before visualization, ensuring the compliance metrics were accurate and meaningful. Challenges included learning a new language that we had never previously used and

ensuring the transformations done to the data were accurate. Our success came through reading through various reference materials online to learn how to write queries in the language.

### **8.2.7 Azure Resources and Policies**

A significant portion of technical learning for both projects involved delving into the intricacies of Azure resources and policies. Understanding how these components interacted and extracting relevant data for the dashboard required in-depth exploration. Challenges included deciphering the complex Azure resource hierarchy and policy structures in order to create a new data model to be used in the dashboard. Diving through a huge number of possible tables in Azure Resource Graph Explorer to retrieve information to use in our data model proved to be a challenge. Success came through a combination of continuous learning and collaboration with various members of the compliance team.

### **8.2.8 Full-Stack Architecture and Technologies**

In order to develop the Compliance History Dashboard project, we needed to learn and build a full-stack architecture, integrating technologies such as React, Vite, and Express to ensure a cohesive and efficient development and deployment process. React, a JavaScript library, serves as the front-end framework, providing a modular and responsive user interface. Vite, chosen as the build tool, enhances the development workflow by supporting the latest ECMAScript features and providing quick development server start-up. Express, a minimal Node.js web application framework, provides a flexible environment for server-side logic, enabling efficient communication between the front-end and other services. In the case of the Compliance History Dashboard, Express manages the logic for uploading PDF compliance reports and facilitating communication between the front-end and Azure Blob Storage.

## 8.2.9 Azure Virtual Machine

Azure Virtual Machine is a fundamental component of cloud computing, providing users with scalable and customizable virtualized computing instances. In hosting environments, Azure VMs offer a secure and controlled space for deploying applications, running backend servers, and ensuring the continuous availability of web services. Leveraging Azure VMs allows for efficient resource management, scalability to meet varying workloads, and a reliable foundation for hosting websites and applications. Hosting the Compliance History Dashboard website on an Azure VM allows for the automated updates of compliance reports to Azure Blob Storage and supports real-time user interaction.

## 8.2.10 Azure Blob Storage

Azure Blob Storage is employed as the storage solution for the generated PDF compliance report files in the Compliance History Dashboard. This cloud-based storage system ensures scalability, security, and accessibility of compliance reports. By utilizing Azure Blob Storage, the project achieves a robust and reliable storage mechanism that facilitates the automated generation and storage of compliance reports at defined intervals.

## 8.3 Accomplishments

The development and implementation of the Governance Compliance Dashboard and the Compliance History Dashboard marked a series of significant accomplishments, showcasing the team's proficiency in both technical and business domains.

The team excelled in creating meaningful visualizations that provided clear insights into compliance metrics. Utilizing Power BI and Power Query, the implementation of Microsoft Query Language (M Query) allowed for sophisticated data preprocessing, enabling the creation of calculated fields, risk scores, and comprehensive metrics. The visualizations, such as compliance pie charts, top non-compliant resources, and policies, significantly enhanced the interpretability of the compliance landscape for the end-users. Furthermore, the team surpassed expectations and successfully completed a second project using React and Azure Blob Storage to showcase potential future features of the dashboard.

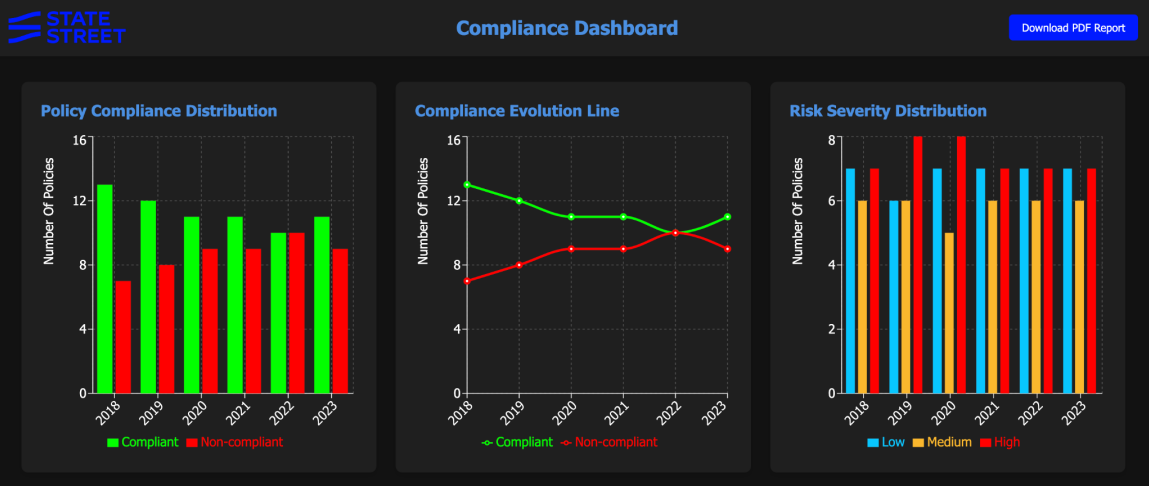


Figure 8.3.1: Compliance History Dashboard

A major accomplishment was the improvement of user experience and accessibility for both the Governance Compliance Dashboard and the Compliance History Dashboard. For the Governance Compliance Dashboard, the decision to create a separate page for specific visualizations, coupled with the maintenance of filters between pages, contributed to a more organized and user-friendly interface. The ability to publish the dashboard on the Power BI website further increased accessibility for the compliance team, allowing them to access and refresh data easily without the need for manual downloads. For the Compliance History Dashboard, an intuitive

interface is demonstrated, as shown in Figure 8.3.1, and we achieved a lighthouse accessibility score of 100, indicating that our dashboard was very accessibility friendly.

The creation of an effective risk assessment model for the Governance Compliance Dashboard was a notable achievement. Integrating KQL queries for Azure Resource Graph Explorer and utilizing M Query Language for data preprocessing, the team developed a comprehensive risk score metric. This metric provided a holistic view of compliance risks, allowing stakeholders to prioritize and address critical issues more effectively.



Figure 8.3.2: Compliance History Dashboard PDF Display

The successful integration of React and Recharts for front-end development for the Compliance History Dashboard, coupled with Vite's efficient build tool, has resulted in a user-friendly interface that seamlessly visualizes historical compliance data. The adoption of Express for the backend server ensures smooth communication between the front-end and Azure services,



facilitating the secure uploading of PDF compliance reports to Azure Blob Storage. The PDFs that are generated will look very similar to Figure 8.3.2, in which the title page displays the State Street logo, a background image of their new headquarters, the title, and the time it was generated. Additionally, the proceeding pages show the policy history data with their compliance status and severity, as well as a watermark to protect confidential information. Leveraging Azure Blob Storage not only ensures the secure storage of compliance reports but also automates their generation at defined intervals, enhancing efficiency and reliability. The use of an Azure Virtual Machine for hosting the website guarantees a scalable and reliable environment, supporting automated updates and real-time user interaction. Most importantly, this project showcases what can be added to the Governance Compliance Dashboard in the future.

Ensuring the sustainability of both projects, the team focused on comprehensive documentation and knowledge transfer. For example, the intricacies of our data model, Graph API queries, M Query transformations, and the overall architecture were documented systematically. This documentation not only served as a guide for future enhancements but also facilitated knowledge transfer within the team, ensuring continuity and scalability.

Collaborating with State Street's international team members, particularly those in India, was a noteworthy accomplishment. The team's contacts in India played integral roles in the projects' success, contributing diverse perspectives and technical expertise. While challenges arose due to time zone differences, accommodations were made, and the collaboration underscored the global nature of the project.

The MQP provided an opportunity to apply theoretical knowledge in a real-world scenario, fostering a deeper understanding of both business and technical aspects. The collaboration with the compliance team exposed the challenges of working with legacy systems and the importance of clear documentation for future enhancements or modifications. Balancing technical intricacies with the

user experience considerations was a valuable lesson, emphasizing the need for effective communication and feedback loops throughout the development process.

In summary, there were many accomplishments, both technical and non-technical, within the two projects, such as automation with data retrieval and integration, compliance report generation, risk assessment modeling, collaboration, and international teamwork. The cumulative success of these accomplishments not only met the initial goals but also positioned the project for sustainability and future enhancements.

## **8.4 Limitations**

### **8.4.1 Change in Project Scope**

The Governance Compliance Dashboard faced various changes in scope, as the team and State Street struggled to define an appropriate project that would both greatly advance our learning and apply all of the knowledge that we have learned during college. Adapting to these changes required agile project management to ensure that goals and deliverables aligned with the adjusted scope. Additionally, starting a new project proved to be extremely difficult due to continuously updating the paper. While the team successfully navigated these modifications, it added an extra layer of challenge in terms of maintaining project coherence and alignment with initial objectives.

### **8.4.2 Limitations in Azure Features for State Street's Cloud**

Despite ambitions to make advancements to the Governance Compliance Dashboard, certain essential Azure features were not enabled for State Street's cloud environment. Restrictions on functionalities like Virtual Machines (VM) and Azure Functions limited the team's ability to explore certain avenues for data processing and application development. This was the primary

reason for completing the second project outside of the State Street environment. This constraint necessitated creative workarounds and prioritization of alternative strategies within the available Azure features.

### **8.4.3 Balancing Technical Advancements and Documentation**

The dual challenge of making technical advancements to the projects while concurrently writing the paper presented a delicate balancing act. The need to document progress, methodologies, and results for the academic paper required significant time and attention. Striking the right balance between hands-on development work and meticulous documentation became essential, and at times, the urgency of paper writing impacted the depth of technical explorations.

## 9. Business and Risk Management

### 9.1 Risks and Costs

In evaluating the risks and costs associated with State Street's new Governance Compliance Dashboard, it is clear that the project's scope inherently limits certain risks. Since the initiative leverages lessons learned from the partially implemented old dashboard, rather than embarking on a completely new development, it sidesteps many hazards typically seen in full-scale development ventures. This strategy capitalizes on the knowledge gained from the prior dashboard implementation, allowing the team to focus on refining and expanding the current functionalities, rather than navigating the unpredictable waters of entirely new development. As a result, this method of leveraging prior experience introduces fewer uncertainties, leading to a more stable and predictable project trajectory. This strategic approach significantly reduces the overall risk profile, making the venture a more solid and dependable undertaking for State Street.

Despite these mitigated risks, there are specific business hazards that need attention during the development of the first project. One prominent concern is the accuracy of data. If the dashboard misrepresents compliance data, State Street and its cloud computing service users might face critical decision-making dilemmas, particularly with resources that are non-compliant. Picture a situation where a customer, misled by the dashboard's data to believe in their environment's compliance, relies on resources that are only temporarily exempt. Once this exemption lapses, their operational environment might abruptly become non-functional.

Another concern with the first project is its omission of a manual refresh feature, a functionality that was available in the old dashboard. In an environment where data is constantly updated, having the option to manually refresh can be extremely valuable. The new dashboard's lack

of this feature restricts users from quickly updating their compliance status, a key factor in making timely and informed decisions. This limitation poses the risk of users making critical choices based on outdated information, as it may delay the visibility of recent changes in compliance data or the status of resources. Such a constraint is a significant step back from the previous system's capabilities, potentially impacting the efficiency and reliability of the decision-making process.

The financial aspects of the first project also warrant a closer look. Operational costs like Azure requests and hosting are on the lower end of the spectrum but are integral to the dashboard's financial landscape. The anticipated maintenance costs are modest, though these might swell if refresh frequencies increase or if automatic data retrieval encounters glitches. A substantial share of the expenses stems from State Street's role in deploying the dashboard. This endeavor involved not only our team but also a significant time commitment from State Street's staff. Our project sponsor and other employees dedicated about 5-10 hours weekly to assist us over a seven-week term, in addition to their involvement during the preceding PQP class. This investment, while not a direct financial cost, represents a considerable allocation of expertise and manpower.

Moreover, the collaboration with the WPI team added its own financial dimensions to the project. While the team itself was not compensated, a nominal fee was paid to their institution, indicative of a direct investment in the development phase. The onboarding of the WPI team brought additional costs, encompassing training, issuance of IDs and State Street laptops, as well as HR and IT support. These elements, though sometimes overlooked, are critical in smoothly integrating external teams into a corporate environment, in addition to labor and other costs on State Street.

Looking ahead, the cost of future upgrades to the first dashboard also needs consideration. Now that State Street will have a functioning dashboard, there will be something tangible to improve and expand upon. This means future upgrade costs will not only be necessary but will also add value

to an already established system, enhancing its capabilities and keeping it aligned with evolving compliance and technological standards. These ongoing improvements are an investment in maintaining the dashboard's relevance and efficiency, ensuring it continues to meet State Street's evolving needs.

In sum, the development of State Street's first compliance dashboard involves a multifaceted investment, balancing mitigated risks with the intricate costs of implementation and future upgrades. This project, significantly aided by the dedicated involvement of State Street personnel and the strategic partnership with WPI, illustrates State Street's commitment to strengthening its compliance management infrastructure, positioning it for future technological advancements.

### **9.1.1 State Street Risk Culture**

In evaluating the risks and costs associated with State Street's new compliance dashboard, it is evident that the project's strategic approach mitigates certain risks, a critical consideration for a systemically important financial institution like State Street. Leveraging lessons from the partially implemented old dashboard instead of embarking on entirely new development minimizes uncertainties, providing a more stable project trajectory. This method reduces the overall risk profile, ensuring the venture is a dependable undertaking for State Street.

Despite mitigated risks, attention must be given to specific business hazards during the dashboard's development. The accuracy of data emerges as a concern, as misleading compliance information could lead to critical decision-making dilemmas for State Street and its cloud computing service users. A particular risk lies in customers relying on temporarily exempt resources, leading to operational challenges once exemptions lapse.

Additionally, the omission of a manual refresh feature in the new dashboard poses a risk. Lacking the option for manual refresh restricts users from promptly updating their compliance

status on-demand. This limitation could result in users making critical decisions based on outdated information, affecting the efficiency and reliability of decision-making processes.

Financial aspects of the project, including operational costs, maintenance expenses, and the investment of State Street's personnel, need careful consideration. The collaboration with the WPI team introduces financial dimensions, encompassing training, onboarding, and support costs. Future upgrades to the dashboard also warrant attention, ensuring ongoing improvements align with State Street's evolving needs.

In summary, the development of the Governance Compliance Dashboard involved a multifaceted investment, balancing mitigated risks with the intricate costs of implementation and future upgrades. This project, significantly aided by the dedicated involvement of State Street personnel and WPI, illustrates State Street's commitment to strengthening its compliance management infrastructure, positioning it for future technological advancements.

### **9.1.2 WPI Team Risk Culture**

The risk culture within the WPI team is characterized by a thorough understanding and proactive management of risks associated with the development of State Street's new compliance dashboard. The team acknowledges the importance of leveraging lessons from the partially implemented old dashboard, opting for a strategic approach that minimizes uncertainties and stabilizes the project trajectory. This risk-aware culture is integral to the team's decision-making process, ensuring that potential pitfalls are identified and addressed in a timely manner.

The team's risk-aware culture is fostered by a combination of factors. Firstly, the team recognizes the significance of prior experience and knowledge gained from the old dashboard implementation. This acknowledgment allows the team to focus on refining and expanding current functionalities rather than navigating the unpredictable challenges of entirely new development.

Secondly, the team prioritizes stability and predictability in the project trajectory, understanding that a strategic approach contributes to a more dependable undertaking for State Street. This risk-conscious mindset is rooted in a commitment to delivering a reliable and effective compliance dashboard.

The team employs various risk controls to ensure the success of the Governance Compliance Dashboard. Leveraging prior experience is a key control measure, providing a foundation of knowledge to navigate potential challenges effectively. The decision to refine and expand current functionalities, rather than engage in entirely new development, acts as a risk mitigation strategy. Regular evaluations of the accuracy of compliance data and the potential impact on decision-making contribute to proactive risk management.

Additionally, the team recognizes specific business hazards, such as the omission of a manual refresh feature, and addresses them head-on. By identifying these risks early in the development process, the team can implement controls to mitigate their impact. The financial aspects of the project are carefully considered, with a focus on operational costs, maintenance expenses, and the investment of State Street's personnel. This financial awareness allows the team to make informed decisions that contribute to the sustainability and success of the project.

Overall, the WPI team's risk culture is characterized by a proactive and informed approach to identifying, assessing, and mitigating risks throughout the development of the compliance dashboard.

## **9.2 Benefits**

The new compliance dashboard at State Street (the first dashboard) brings multiple business benefits, contributing to the company's operational efficiency, customer satisfaction, and strategic development.



### **9.2.1 Time Efficiency and Operational Productivity**

The new dashboard offers a notable advantage in terms of saving time. Previously, the Azure operations team at State Street had to dedicate approximately 2-3 hours per week to manually refresh data. Managing compliance data manually involved the tedious task of updating Excel sheets regularly, demanding constant attention and updates. In the old system, this process was further hindered by lengthy refresh times, where employees often waited over half an hour for PowerBI to update after manually inputting data into Excel sheets. This led to delays and inefficiencies in data management and analysis. With the introduction of the new dashboard, this process has become automated, eliminating the lengthy wait times and streamlining data updates. As a result, there is a significant time saving for the Azure Operations team. This time efficiency translates into improved productivity, enabling the team to shift their focus towards more strategic initiatives and proactive problem-solving, rather than being tied down by repetitive administrative tasks. Under the old system, routine updates and manual data management incurred significant expenses. By automating these processes, the new dashboard minimizes the need for continuous human intervention, reducing the hours dedicated to maintenance and consequently lowering the overall maintenance budget. This cost-effectiveness is not just a short-term benefit but also a long-term advantage, contributing to the sustainability of the project.

### **9.2.2 Enhancing Tenant Experience through Self-Service**

The self-service aspect of the new dashboard is a major leap forward in customer experience. Previously, customers reliant on State Street's cloud services had limited access to compliance information. They had to request reports and rely on State Street personnel to provide them, or in some cases, even write their own visualizations. This process was not only time-consuming for the customers but also created a bottleneck in terms of efficiency for State

Street employees. The lack of a functional dashboard led to significant frustration, as employees had to manually find and compile data without the help of the dashboard, interfering with their primary work responsibilities and leading to slower service delivery, which in turn frustrated customers.

### **9.2.3 Scalability and Client Management**

As State Street's clientele continues to expand, the importance of the scalability of its services becomes increasingly evident. The dashboard plays a pivotal role in achieving this scalability. In the past, managing a substantial number of clients was manageable but lacked the potential for scalability. As the client base grows, the manual approach becomes progressively unwieldy and inefficient, necessitating team expansion to handle the increasing workload manually.

The new dashboard, equipped with self-service and automated functionalities, empowers State Street to effectively oversee a larger volume of clients. This scalability is paramount for business growth, as it ensures that the quality of service remains uncompromised despite the expanding client base. For example, if State Street were to extend its client roster to hundreds, the self-service capabilities of the dashboard would obviate the need for a proportional increase in the operations team size. This scalability not only enhances operational efficiency but also paves the way for business expansion without incurring corresponding increases in operational expenditures.

### **9.2.4 Security and Compliance Monitoring**

The dashboard has the potential to significantly enhance the monitoring of security and compliance by introducing automated features. In the financial sector, strict compliance standards are not only a regulatory requirement but also a fundamental element of customer trust and business integrity. The new dashboard aims to provide a centralized platform for monitoring compliance and security standards, with the possibility of introducing automated functions in the future. Platform

managers and app owners will have access to a streamlined, updated view of their compliance status, enabling quicker responses to compliance shortfalls. For instance, if a particular tenant's compliance status were to drop, the dashboard could be extended to flag this issue automatically, allowing for prompt investigation and resolution by State Street. Proactive monitoring is essential for maintaining high standards of security and compliance, ultimately reinforcing State Street's reputation as a secure and reliable financial services provider.

### **9.2.5 Access Control and Data Segregation**

A key feature of the new dashboard is its enhanced access control and data segregation capabilities. In the old system, access to compliance data was limited, often confined to a few individuals with portal access. This restriction not only created bottlenecks in data access but also raised concerns about data privacy and segregation. The new dashboard addresses these issues by implementing robust access controls, ensuring that users can only view data relevant to their specific roles and responsibilities. This feature is crucial in maintaining the confidentiality of sensitive information across different clients and departments. For instance, a tenant manager can access only the compliance data pertinent to their managed assets, preventing any unauthorized viewing of other tenants' data. Such precise control over data access not only enhances security but also ensures compliance with stringent regulatory standards regarding data privacy and handling.

### **9.2.6 Long-Term Development and Strategic Insights**

The new dashboard is not just a solution for current needs but also a foundation for future technological advancements. Its design allows for future integrations, such as auto-remediation capabilities, which could further streamline compliance management processes. For example, in the future, the dashboard could offer functionality to automatically correct certain compliance issues

with minimal human intervention, thereby enhancing efficiency and reducing the risk of human error. Moreover, for higher-level managers and decision-makers, the dashboard provides valuable insights into the platform's overall performance and risk profile. They can track improvements over time, assess compliance trends, and make informed strategic decisions. For instance, platform managers can use the dashboard to monitor the overall compliance rate across different tenants, identifying areas where guidelines are not being strictly followed and taking corrective actions. This level of oversight is vital for effective governance and strategic planning.

### **9.2.7 Monitoring and Governance**

The dashboard plays a pivotal role in monitoring adherence to State Street's compliance guidelines. It provides comprehensive visibility into how many tenants are following the prescribed standards, which is essential for consistent governance across the platform. This monitoring capability enables State Street to identify and address potential risks before they escalate. For example, if a certain application shows a pattern of non-compliance, the dashboard can quickly bring this to the attention of the relevant managers, prompting immediate action.

### **9.2.8 Impact on Customer Confidence and Business Growth**

An important aspect of the dashboard is its potential impact on customer confidence and, consequently, on business growth. By transparently displaying improvements in security and compliance metrics, the dashboard can significantly boost client trust. For instance, a visible improvement in compliance metrics from 60% to 95% can greatly enhance customer confidence, potentially leading to increased business and client retention.

The potential impact of the dashboard on customer confidence and, subsequently, on State Street's business growth, cannot be overstated. In the financial services industry, where trust and

security are paramount, the dashboard's ability to transparently display security and compliance metrics plays a crucial role in strengthening client relationships.

The dashboard's transparent display of compliance metrics serves as tangible proof of State Street's commitment to maintaining high standards in security and compliance. For instance, a marked improvement in these metrics, such as an increase from 60% to 95% compliance, is not just a numerical change but a significant indicator of State Street's diligence and reliability. Such visible enhancements in compliance can profoundly bolster customer confidence. Clients, when able to see updated data and improvements in their compliance metrics, are reassured about the safety and regulatory adherence of their assets managed by State Street. This level of transparency and accountability is a powerful tool in building and maintaining trust. In an industry that is highly competitive and driven by trust, the dashboard offers State Street a distinctive edge. Prospective clients, when comparing service providers, are likely to be influenced by the demonstrable compliance standards and the technological sophistication that State Street offers. The dashboard can become a key factor in their decision-making process. Clients looking for a financial services cloud provider will find the visible commitment to compliance and security a compelling reason to choose State Street over other providers. This advantage is especially significant in acquiring new clients who are navigating the complex landscape of financial services and seeking a provider they can trust with their assets and data.

For existing clients, the dashboard reinforces their decision to partner with State Street. It offers them a continuous insight into how their assets are being managed, with a focus on compliance and security. This ongoing assurance plays a vital role in client retention. In the financial sector, where switching service providers can be a complex and risky endeavor, clients are more likely to stay with a provider that consistently demonstrates its capability and commitment to maintaining high standards of compliance and security.

The improved customer confidence directly translates into business growth. Satisfied clients are more likely to increase their investments and use a wider range of services offered by State Street. Additionally, high customer satisfaction and trust often lead to word-of-mouth referrals, attracting new clients and further expanding the business. The reputation of a financial institution is one of its most valuable assets, and the dashboard significantly contributes to enhancing State Street's reputation in the market.

The dashboard's impact on customer confidence also aligns with State Street's long-term strategic goals. In an ever-evolving financial landscape, where regulatory requirements and security challenges are becoming increasingly complex, State Street positions itself as a forward-thinking, technologically adept company. This positioning not only helps in current client satisfaction and acquisition but also lays the groundwork for future innovations and adaptations in the rapidly changing financial sector.

The new compliance dashboard at State Street is a multifaceted tool that brings significant business value. It enhances operational efficiency, reduces costs, improves customer experience, and provides strategic insights. The dashboard's scalability and potential for future enhancements position State Street to continue its growth while maintaining high standards of compliance and security, pivotal in the financial services sector.

### **9.3 Compliance History Dashboard**

The development of the second dashboard at State Street, envisioned as a proof of concept, represents an exploratory step in the evolving landscape of financial services data management and governance. As an independent React application, it was designed to provide a conceptual model for detailed visualizations of historical compliance data. However, while the dashboard takes strides in

conceptual compliance management, it does not utilize actual State Street data or resources, rendering it a non-operational prototype in its current state.

The primary purpose of this proof of concept is to demonstrate the potential future additions to the Governance Compliance Dashboard, such as compliance history visualizations. By integrating Recharts, the dashboard mock-up is designed to showcase how historical compliance data can be presented in various formats, like bar charts and line charts. The intention is to illustrate how such a tool could enable organizations like State Street to track compliance trends over extended periods. These capabilities are critical for understanding long-term trends and patterns in compliance, offering theoretical insights that could be valuable for strategic planning and forecasting future compliance needs. However, since this prototype operates without real State Street data, its current utility in strategic decision-making and policy assessment is hypothetical. This mock-up also highlights the potential for streamlining the auditing and record-keeping processes. Ideally, such a dashboard could serve as an organized and easily accessible repository of historical compliance data, significantly simplifying audit processes. However, as the prototype does not incorporate actual data, its practical application in internal audits and regulatory reviews remains a concept to be realized in future implementations.

The development of this proof of concept required dedicated time and resources from the WPI team, despite its current non-operational status. The use of technologies like React and Recharts in this mock-up required specific technical expertise. While the development of a fully functional dashboard using these technologies might extend the timeline and add to the complexity of the project, the current prototype serves as a valuable learning experience and a stepping stone for future development.

While the second dashboard at State Street stands as a promising conceptual tool for enhancing compliance management, its status as a mock-up without integration of actual data or

resources from State Street means it currently serves more as an illustrative model. Its development underscores the potential benefits and technical requirements of such a tool, providing a foundation for future projects that could fully realize these capabilities with the integration of real data and resources.



## 10. Future Work

Future work to build on the foundation of the Governance Compliance Dashboard and Compliance History Dashboard projects at State Street has the potential for expansive growth and enhancement. The following outlines potential avenues for future work and development of both projects:

### 10.1 Integration of Additional Azure Features

Expanding the project to leverage currently disabled Azure features on the State Street cloud environment, such as Virtual Machines (VM) and Function Apps, would significantly enhance the dashboard's capabilities. This could involve developing more sophisticated data processing algorithms, implementing advanced automation, and incorporating serverless functions for dynamic responses to compliance changes.

### 10.2 Integration of the Second Dashboard

Alongside section 10.1, the integration of the second dashboard, which currently stands as a proof of concept for historical data visualization, with the existing first dashboard at State Street, presents an exciting avenue for future development. This integration could create a more comprehensive and cohesive compliance management system as additional Azure features are enabled on State Street's cloud environment, such as Azure Virtual Machines. By blending the capabilities of both dashboards, State Street could benefit from existing compliance monitoring (offered by the first dashboard) alongside in-depth historical compliance analysis (provided by the second dashboard). Such integration would involve synchronizing data flows and user interfaces

between the two dashboards, ensuring seamless navigation and data consistency. This could allow users to switch effortlessly between viewing current compliance statuses and historical compliance trends, providing a more rounded perspective on their governance framework. This unified system could also facilitate more informed decision-making by allowing users to correlate current compliance data with historical patterns.

### **10.3 Machine Learning for Predictive Analytics**

Integrating machine learning algorithms for predictive analytics could be a transformative addition to the Governance Compliance Dashboard. With ample resources, the team could explore and implement models that predict potential compliance risks based on historical data, providing proactive insights to the compliance team.

### **10.4 Enhanced User Interactivity and Personalization**

Investing in the development of a more interactive and personalized user experience could be a valuable avenue. This might involve incorporating advanced data visualization techniques, allowing users to customize their dashboards based on individual preferences and specific compliance metrics of interest.

### **10.5 Integration with External Risk Management Systems**

Considering integration with external risk management systems or threat intelligence platforms could further enhance the dashboard's capabilities. This could involve real-time synchronization with external databases, providing a broader context for compliance metrics and enriching the analysis with external risk factors.

## 10.6 Automated Remediation

Integrating automated remediation strategies would be a logical extension of the project. In the future, the team could design and implement automated workflows to address non-compliant resources in real-time. This would involve creating predefined remediation actions based on compliance policies, reducing the manual effort required to rectify issues, and enhancing the agility of the governance framework.

## 11. Conclusion

We embarked on two ambitious tasks: developing a compliance visualization dashboard to enhance cloud governance in the financial sector and creating a second, conceptual dashboard to build off of the first one for State Street. These ventures extended beyond mere technical challenges, presenting us with the opportunity to address significant operational needs within State Street and offering a realistic perspective on the practical applications of our academic knowledge.

Throughout this endeavor, our team achieved significant milestones in both projects. For the Governance Compliance Dashboard, we excelled in creating meaningful visualizations using Power BI and Power Query, yielding clear insights into compliance metrics. This was evident in our compliance pie charts, top non-compliant resources, and policies, enhancing end-user interpretability. In contrast, the Compliance History Dashboard, a proof of concept, was a non-operational mock-up designed as an independent React application. It provided a conceptual model for visualizing historical compliance data, highlighting the potential of dynamic visualization for tracking compliance trends and streamlining auditing processes.

A major accomplishment across both projects was the improvement of user experience and accessibility. We created separate pages for specific visualizations and incorporated filters, delivering an organized and user-friendly interface. Additionally, in the compliance visualization project, we developed an effective risk assessment model by integrating KQL queries and M Query Language, offering a holistic view of compliance risks. The sustainability of both projects was also a key focus. We ensured comprehensive documentation and knowledge transfer for continuity and scalability. While the first project's development involved sophisticated data preprocessing and risk assessment models, the second project's development underscored the technical requirements and potential

benefits of such tools, setting the stage for future projects that could integrate real data and resources.

Throughout both projects, we encountered challenges that tested our problem-solving skills and adaptability. Integrating compliance data into a user-friendly interface required deep engagement with technical aspects of cloud computing and the regulatory landscape of the financial sector. Despite these hurdles, our team demonstrated resilience and ingenuity, effectively navigating these complexities using the Agile Scrum methodology.

Our collaboration with State Street has been a cornerstone of our learning experience, enabling us to make a tangible impact in the real world. The insights and experiences gained from both projects are invaluable, extending well beyond our academic journey. We are confident that the findings and dashboards developed in these projects will not only meet the initial goals but also set the stage for future enhancements, charting a new path in the realm of cloud governance.

## References

Amazon Web Services. (2023). *What is cloud computing?*

<https://aws.amazon.com/what-is-cloud-computing/>

Atlassian. (2023). *Jira Software - Features*. <https://www.atlassian.com/software/jira/features>

Beck, K.L., Beedle, M.A., Bennekum, A.V., Cockburn, A., Cunningham, W., Fowler, M., Grenning, J., Highsmith, J., Hunt, A., Jeffries, R., Kern, J., Marick, B., Martin, R.C., Mellor, S.J., Schwaber, K., Sutherland, J., & Thomas, D.A. (2013). *Manifesto for Agile Software Development*.

<https://agilemanifesto.org>

Caceres, Z. (2022). Express.js – A gentle overview. Express.js – A Gentle Overview.

<https://gist.github.com/zcaceres/23d9d6eb6eb74947bb05ab4d3c996499>

Celner, A. (2019, July 1). *Cloud banking: Financial services and banking of the future*. Deloitte.

<https://www.deloitte.com/global/en/Industries/financial-services/perspectives/bank-2030-financial-services-cloud.html>

Davidsmatlak. (2023a, June 15). *Overview of azure policy - azure policy*. Overview of Azure Policy - Azure Policy | Microsoft Learn.

<https://learn.microsoft.com/en-us/azure/governance/policy/overview>

Davidsmatlak. (2023b, October 31). *Overview of azure resource graph - azure resource graph*. Overview of Azure Resource Graph - Azure Resource Graph | Microsoft Learn.

<https://learn.microsoft.com/en-us/azure/governance/resource-graph/overview>

Dubey, A. (2023, October 10). Introduction to Azure Blob Storage. Introduction to Blob (object) Storage - Azure Storage | Microsoft Learn.

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blobs-introduction>

EidosMedia. (2022, April 12). *The Future of Cloud Computing in Fintech. Learn why cloud computing in fintech is essential.* <https://www.eidosmedia.com/blog/tecnology/cloud-computing-in-fintech>

History, C. (2023, October 30). *State Street Corporation.* CompaniesHistory.com - The largest companies and brands in the world.

<https://www.companieshistory.com/state-street-corporation/>

IBM. (2023). *What is cloud computing?* <https://www.ibm.com/topics/cloud-computing>

James, M. (2023). *Scrum Reference Card.* <https://scrumreferencecard.com/scrum-reference-card/>

JnHs, et al. (2022, September 13). *What Is the Azure Portal? Azure Portal Documentation.*

<https://learn.microsoft.com/en-us/azure/azure-portal/azure-portal-overview>

Liberto, D. (2023, March 27). *Systemically Important Financial Institution (SIFI) Overview.* Investopedia.

<https://www.investopedia.com/terms/s/systemically-important-financial-institution-sifi.asp>

Librarian. (2023). *Beginner's Guide to Discord.* Discord Basics.

<https://support.discord.com/hc/en-us/articles/360045138571-Beginner-s-Guide-to-Discord>

Lopez, N., et al. (2022, December 13). *Chaos Engineering for Cloud Resiliency.* Digital WPI.

[https://digital.wpi.edu/concern/student\\_works/pz50h0762?locale=en](https://digital.wpi.edu/concern/student_works/pz50h0762?locale=en)

McIntosh, N. (2021, March 8). *The impact of cloud computing in Fintech.* Telehouse.

<https://www.telehouse.net/blog/the-impact-of-cloud-computing-in-fintech/>

Meta Open Source. (2023). Quick start. React. <https://react.dev/learn>

Microsoft. (2023a). Power BI. Power BI - Data Visualization | Microsoft Power Platform.

<https://www.microsoft.com/en-us/power-platform/products/power-bi/>

Microsoft. (2023b). *What is cloud computing?: Microsoft Azure*. What Is Cloud Computing? | Microsoft Azure.

<https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-is-cloud-computing>

Neugebauer, N. (2023, September 29). What is Azure SQL Managed Instance?. Azure SQL Managed Instance | Microsoft Learn.

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview?view=azuresql>

Nottingham, C. (2023, February 28). Virtual machines in Azure. Overview of virtual machines in Azure - Azure Virtual Machines | Microsoft Learn.

<https://learn.microsoft.com/en-us/azure/virtual-machines/overview>

Pattnaik, A. (2023, October 5). *Announcing General Availability of the New Microsoft Teams App for Windows and Mac*. Tech Community Microsoft.

<https://techcommunity.microsoft.com/t5/microsoft-teams-blog/announcing-general-availability-of-the-new-microsoft-teams-app/ba-p/3934603>

Recharts Group. (2023). Recharts. <https://recharts.org/>



Rosenberg, E. (2022, November 27). *The Importance of Excel in Business*. Investopedia.

<https://www.investopedia.com/articles/personal-finance/032415/importance-excel-business.asp>

Rubin, K. S. (2013). *Essential Scrum: A Practical Guide to the Most Popular Agile Process*. Addison-Wesley.

SailPoint Technologies, Inc. (2021, April 28). *The importance of a cloud governance framework*. SailPoint.

<https://www.sailpoint.com/identity-library/cloud-governance-framework/>

Shacklett, M. (2023, November 8). *What is cloud compliance?: A comprehensive guide*. Datamation.

<https://www.datamation.com/cloud/what-is-cloud-compliance/>

Shsagir, et al. (2023, August 27). *Kusto Query Language (KQL) Overview*. Azure Data Explorer & Real-Time Analytics | Microsoft Learn.

<https://learn.microsoft.com/en-us/azure/data-explorer/kusto/query/>

State Street. (2023). *Meet Our Company*.

<https://www.statestreet.com/us/en/asset-manager/about/our-story>

State Street. (2023, January 26). *State Street Announces Strategic Cloud and Infrastructure Solution Providers in Connection with Its Multi-Year Technology Transformation Journey*.

<https://investors.statestreet.com/investor-news-events/press-releases/news-details/2023/State-Street-Announces-Strategic-Cloud-and-Infrastructure-Solution-Providers-in-Connection-with-its-Multi-Year-Technology-Transformation-Journey/default.aspx>

Sudhir, S. (2022, October 10). *What is Azure Automation?*. Azure Automation overview | Microsoft Learn. <https://learn.microsoft.com/en-us/azure/automation/overview>

Suñé, J. (2023, February 6). *Continuous Cloud Optimization Insights (CCO Insights) - Home*. CCOInsights.

<https://github.com/azure/CCOInsights/wiki>

Wang, B., & Citrix Staff. (2023, June 8). *Technical Overview*. Technical Overview | Citrix Virtual Apps and Desktops 7 2308.

<https://docs.citrix.com/en-us/citrix-virtual-apps-desktops/technical-overview>

You, E. (2023). Vite. <https://vitejs.dev/guide/>

## Appendix A: Glossary of Terms

**Agile Scrum Methodology:** A framework for managing and completing complex projects, known for its adaptability and iterative process.

**Azure DevOps:** A suite of development tools provided by Microsoft, designed to support collaborative software development projects. Azure DevOps includes features for version control, reporting, requirements management, project management, automated builds, testing, and release management, facilitating a comprehensive and integrated approach to software development and deployment.

**Azure Services:** Microsoft's cloud computing services, including various tools and platforms for managing cloud resources.

**Azure Policy:** A service in Azure for enforcing organizational standards and assessing compliance.

**Azure Portal:** A web-based interface for managing Azure subscriptions and services.

**Azure Resource Graph:** A service for exploring Azure resources and their properties.

**Citrix Workspace:** A digital workspace platform that integrates various applications and services.

**Cloud Computing:** The delivery of computing services over the Internet, including servers, storage, databases, networking, software, analytics, and intelligence.

**Cloud Governance:** The process of overseeing and controlling cloud computing products and services.

**Cloud Migration:** The process of moving data, applications, or other business elements to a cloud computing environment.

**Compliance/Governance Dashboard:** A tool for visualizing and managing compliance data in cloud governance.

**Compliance Monitoring:** The process of ensuring adherence to established regulations, standards, and internal policies.

**Data Visualization:** The graphical representation of information and data to provide an accessible way to see and understand trends, outliers, and patterns.

**Data Security:** Measures taken to protect digital information from unauthorized access or alteration.

**DevOps:** A set of practices that combines software development (Dev) and IT operations (Ops) to shorten the systems development life cycle and provide continuous delivery.

**Entity-Relationship Diagram (ERD):** A graphical representation used in database design to illustrate the relationships between different data entities in a system. ERDs depict entities (such as tables in a database) and how they interrelate, showing relationships, primary and foreign keys, and other database schema elements. They are crucial for visualizing the structure of database systems.

**Financial Technology (Fintech):** Technology and innovation aimed at competing with traditional financial methods in the delivery of financial services.

**Infrastructure as a Service (IaaS):** A form of cloud computing that provides virtualized computing resources over the internet.

**Kusto Query Language (KQL):** A language used for querying Azure data services.

**Machine Learning:** A type of artificial intelligence that allows software applications to become more accurate at predicting outcomes without being explicitly programmed to do so.

**Microsoft Excel:** A spreadsheet program used for data entry, analysis, and visualization.

**Microsoft Graph API:** An API for accessing data from Microsoft 365 services.

**Minimum Viable Product (MVP):** A development approach focusing on creating a basic version of a product with essential features to satisfy early users and gather feedback for future enhancements.

**Power BI:** A business analytics service by Microsoft for data visualization.

**Power Query:** A data connection technology that enables you to discover, connect, combine, and refine data sources to meet your analysis needs.

**Predictive Analytics:** The use of data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes based on historical data.

**React:** A JavaScript library for building user interfaces, notable for its efficient and dynamic rendering of web components. Developed by Facebook, React employs a declarative JSX syntax and component-based architecture, making it popular for creating interactive and scalable single-page applications.

**Regulatory Compliance:** Adherence to laws, regulations, guidelines, and specifications relevant to business processes.

**Risk Management:** The process of identifying, assessing, and controlling threats to an organization's capital and earnings.

**SQL Database:** A structured query language database, used for managing and querying relational databases.

**State Street Corporation:** A financial services and bank holding company.

**Systemic Important Financial Institutions (SIFIs):** A bank, insurance company, or other financial institution whose failure might trigger a financial crisis.

**Virtual Desktop Infrastructure (VDI):** Technology that hosts a desktop operating system on a centralized server in a data center.