disc
Data Integration Support Center
WestEd

2025
June

SPOTLIGHT

Maintaining a Secure AI Environment

RESOURCE CONSIDERATIONS

AUTHORS:
Baron Rodriguez
Laia Tiderman
Sara Kock

Public agencies are increasingly turning to generative artificial intelligence (GenAI) to gain actionable insights from their data, streamline operations, and create innovative solutions that benefit their citizens. The benefits can be compelling, but they are not instant—creating a safe, secure AI environment requires careful planning, thoughtful budgeting, and steady commitment.

This brief from the Data Integration Support Center (DISC) at WestEd examines the resources needed to set up and maintain a secure AI environment. The information is informed by WestEd's experience implementing its own secure internal AI environment in response to staff members increasingly requesting access to GenAI tools to enhance efficiency.

The brief includes examples related to personnel, software and licensing, cloud environment, and ongoing operations. Understanding these practical considerations is essential for agencies to make informed decisions about their AI implementation strategy and to ensure that their investment creates lasting value while maintaining security and trust. A companion brief, "Building a Secure AI Environment: Implementation Plan," outlines the steps involved in implementation and provides details on WestEd's use case.

Before adopting AI capabilities, organizations should consider several factors. DISC strongly recommends having a clear understanding of the risks, governance, benefits, challenges, staffing, training, and specific use cases for the organization. Many of these topics and how WestEd worked on them are outlined in prior briefs. Additionally, DISC offers support to public agencies in facilitating discussions and forming strategies to create a strong foundation for AI implementation.

Resource Considerations

The resources associated with a secure AI environment fall into four major categories: personnel, software and licensing, cloud environment, and ongoing operations. Each category represents distinct investments that organizations need to consider in their planning and maintenance.

- **Personnel**
- **Software and Licensing**
- **Cloud Environment**
- **Ongoing Operations**

PERSONNEL

User Community

Staff members outside of the technology department play an important role. These individuals participate in use case development, trainings, and testing. The organization needs to consider these efforts and budget accordingly. In addition, agencies need to plan how to support the user community. Support may include an operations team with the following members:

- **training specialists** to create and design content
- **a change management specialist** to monitor and approve environmental modifications (e.g., introduction of new AI models or version updates)
- **support specialists** for user requests and assistance

Technical Team

WestEd's secure AI environment required two primary personnel to oversee the design and technical implementation of the secure AI framework. While titles may differ, the implementation required skill sets consistent with the following positions:

- **a cloud administrator or cloud architect**
- **an information technology administrator** with a specialty in AI and machine learning operations

These staff members may bring in other team members with skill sets consistent with the following positions:

- **an AI engineer** with experience in the implemented model who can test and validate AI models

- **data analysts/engineers** who can work on structuring and classifying data
- **an AI ops engineer** who can automate and deploy the AI models
- **a Python developer** to customize, enhance, and tune the AI models

Additional technical staff will be needed to specialize in or oversee key components of the environment. Those positions require skill sets consistent with the following positions:

- system administrators and security teams, including
  - **a network security engineer** for segmentation of the environment and security groups
  - **a systems administrator** to oversee the virtual desktop infrastructure setup
  - **a security engineer or analyst** to oversee security testing and monitoring
  - **an identity management specialist** to assist with authentication and authorization
- security operations and compliance teams, including
  - **an information security officer** who oversees the overall security strategy for the organization
  - **a security operations analyst** to identify and mitigate potential security risks, monitor alerts and incidents, and perform formal security audits
  - **a compliance officer** to confirm that the regulatory requirements are being met

Even after the secure AI environment is implemented, many of these technical staff will still be needed for ongoing support of the system. However, their level of effort will likely drop as they shift to maintenance mode.

SOFTWARE AND LICENSING

Implementing a secure AI environment requires various software components and associated licenses, each serving specific functions in the overall architecture. Organizations must carefully consider both initial licensing costs and ongoing subscription fees for these essential components. Organizations should also carefully evaluate the different licensing models available for the software (e.g., per user vs. enterprisewide).

Core AI Platform

The foundation of the environment requires licensing for AI development and deployment platforms. Licensing includes costs for enterprise AI services, model development tools, and associated APIs. Organizations should consider both development environments for training models and production environments for deployment.

Virtual Desktop Software

Virtual desktop implementation requires several key software components. These components include operating system licensing for each virtual desktop instance, productivity software subscriptions for user applications, and management tools for the virtual desktop environment. For example, operating systems may rely on Windows 10 Enterprise or Windows 7 Enterprise with Extended Security Updates. Productivity software may include Microsoft 365 subscriptions and statistical software such as R or Stata.

Other Software

Organizations typically need additional software licenses to ensure comprehensive security and efficient management of their secure AI environment. These additional licenses include identity and access management software, security monitoring tools, and system management platforms. For example, organizations may need licenses to implement Azure

Security Center and Azure Sentinel for threat detection and monitoring, Azure Key Vault for key management, Azure Active Directory for access control, and Azure Firewall for network security. Similarly, to enforce access controls, some organizations may need to integrate their secure AI environment into their single sign-on solution. Costs will depend on whether the identification management is already in the cloud or on premises as well as any additional features such as multifactor authentication or single sign-on integration.

Organizations should review their existing enterprise software agreements to determine if current licenses can be extended to cover the AI environment or if new licenses are required.

CLOUD ENVIRONMENT

A secure AI environment requires robust cloud infrastructure that supports both AI operations and secure user access. Organizations must plan for several critical infrastructure components that work together to create a secure, scalable environment.

The plans should include detailed cost estimates based on organizational needs and infrastructure choices. Most cloud providers offer pricing calculators and documentation that can help estimate costs based on specific requirements. While costs vary by provider and implementation decisions, understanding the major cost components helps organizations plan their investments effectively.

Compute Resources

The secure AI environment requires various types of compute resources to support different functions.

- **Virtual machines (VMs):** The selected cloud provider will need VMs to host the virtual desktop infrastructure. The cost of VMs depends on the type, size, and number required to support the expected user base and workload demands. Cloud providers offer various VM families optimized for different use cases. For example, Microsoft Azure offers options such as general purpose, memory optimized, or GPU instances. Cloud providers also have different pricing tiers. Organizations will need to determine which pricing tier is best for them by balancing cost with performance and capabilities.
- **AI model:** Compute resources are required to support the AI tool itself, especially for model training, fine-tuning, and deployment. This support may require specialized instances optimized for machine learning workloads.

Regardless of the current needs, organizations should plan for both baseline capacity and the ability to scale resources as needed.

Storage Infrastructure

Organizations need to account for costs associated with storing and managing large amounts of data in their secure AI environment. Requirements include primary storage for user profiles and applications, data storage for the AI training and operations, backup storage for disaster recovery, and archive storage for compliance.

Cloud storage costs typically follow a consumption-based model in which organizations pay for what they use—and AI consumes a lot of storage. The costs are influenced by several factors: storage type (such as hot storage for frequently accessed data vs. cool storage for infrequently accessed data), data volume, number of storage transactions, and data movement between cloud regions or to/from on-premises systems. For example, frequently accessed AI training data might require hot storage with higher costs, while archived model versions could use cool storage at lower rates. Organizations should also consider that moving data in and out of cloud storage often incurs additional charges. Planning data storage location and access patterns carefully can help manage these costs effectively.

Other Components

The cloud environment requires several critical networking and security components:

- network security groups and firewalls for traffic control
- load balancers for traffic management
- virtual networks for environmental isolation
- identity and access management systems
- security monitoring and threat detection tools
- encryption management systems

These infrastructure components each carry associated costs. Organizations should consider both the initial setup costs and ongoing operational expenses for each component. Network costs require particular attention, as they include both data transfer charges between services and egress charges for data leaving the cloud environment. Organizations should carefully plan their network architecture to optimize these costs while maintaining necessary performance and security levels.

ONGOING OPERATIONS

Maintaining a secure AI environment requires continuous effort and resources beyond the initial implementation. Organizations considering similar implementations should carefully assess their current capabilities, available resources, and long-term commitment to supporting and evolving their AI environment. Maintaining a secure AI environment is an ongoing commitment that requires dedicated resources and continuous attention. Several ongoing operational components are necessary to ensure that the environment remains secure, efficient, and valuable to users.

System Maintenance and Updates

Regular maintenance is essential for both security and performance. This maintenance includes the following:

- software patches and updates across all components
- AI model updates, patches, and refinements
- virtual desktop environment maintenance
- security tool updates and configuration management

Organizations should establish clear update schedules and maintenance windows that minimize disruption to users while ensuring timely security updates.

Monitoring

A comprehensive monitoring strategy is needed for maintaining both security and performance of the secure AI environment. Organizations must implement real-time security monitoring and threat detection systems that continuously scan for potential security issues and unauthorized access attempts. This monitoring should be paired with regular security assessments and vulnerability scanning to proactively identify potential weaknesses before they can be exploited.

Performance monitoring goes hand in hand with security oversight, as system performance can affect both security and user experience. Organizations need to regularly monitor resource usage across the environment, analyzing patterns to identify potential bottlenecks or areas for optimization. This work includes tracking compute resource use, storage consumption, and network performance to ensure that the environment operates efficiently and scales appropriately to meet user needs.

Token Monitoring and Usage Analytics

Tokens are the basic units that AI language models use to process text, representing words, parts of words, or punctuation marks. Most AI services charge based on token consumption, billing separately for input tokens sent to the model and output tokens generated by the model. Organizations must implement token monitoring

systems to track usage across users, applications, and time periods to maintain cost control and operational visibility. This monitoring enables identification of resource-intensive applications, peak usage patterns, and the comparative efficiency of different prompting strategies. Token monitoring also serves as a security control, as unusual consumption patterns can indicate unauthorized access, system misuse, or potential security breaches that warrant investigation.

Without proper token monitoring, AI usage costs can escalate rapidly and unpredictably, potentially creating significant budget overruns. Comprehensive token monitoring and usage analytics provide organizations with the data necessary to forecast future resource needs, establish accurate budget projections, and identify specific opportunities to optimize efficiency and reduce operational costs. These insights are essential for sustainable GenAI implementation and long-term cost management.

Compliance

Compliance requirements add another layer of complexity to monitoring efforts. Organizations must maintain detailed audit trails and documentation to demonstrate adherence to relevant regulations and internal policies. This documentation includes regular compliance audits, access control reviews, and comprehensive logs of system activities. The compliance program should also include periodic reviews of security configurations and policies to ensure that they remain aligned with current requirements and best practices.

Cost optimization should be an ongoing focus of monitoring activities. Regular reviews of resource usage patterns can identify opportunities to optimize costs without compromising security or performance. Optimization might involve adjusting resource allocation, reviewing storage tiers, or fine-tuning network configurations based on actual usage patterns.

Final Thoughts

The implementation of WestEd's secure AI environment provides a concrete example of the staff time, technical capabilities, and financial investment needed for establishing an enterprise AI solution. While every organization's journey will differ based on factors such as size, existing infrastructure, security requirements, and desired AI capabilities, WestEd's experience offers valuable insights into the scope and scale of such an implementation.

The level of effort for implementing WestEd's secure AI environment was significant. It required dedicated resources with expertise in cloud services, AI integration, and cybersecurity. While the initial implementation timeline spanned about 3 months, success depended not only on technical execution but also on strong governance, clear communication channels, and organizational buy-in. WestEd enhanced project success by engaging outside experts in specialized areas throughout the implementation process.

Organizations considering similar implementations should carefully assess their current capabilities, available resources, and long-term commitment to supporting and evolving their AI environment. Maintaining a secure AI environment is an ongoing commitment that requires dedicated resources and continuous attention. Regular evaluation of the environment's effectiveness, security posture, and alignment with organizational needs ensures that the investment will continue to deliver value while maintaining necessary security controls.

How DISC Can Help

DISC at WestEd can facilitate productive conversations and assist in the development of an AI strategy for an organization's integrated data system. DISC offers technical assistance to public agencies free of cost, including the following services:

- training agency leadership and/or staff in GenAI best practices
- providing expertise on the scope and variety of AI tools to support data integration efforts
- developing use cases that leverage AI tools and aligning those use cases to the state's strategic priorities
- conducting neutral, external assessments of the maturity of the integrated data system and its capacity to support and scale AI technologies
- facilitating structured discussions to develop a foundation for the use of AI in the integrated data system

Suggested citation: Rodriguez, B., Tiderman, L., & Kock, S. (2025, June). *Maintaining a secure AI environment: Resource considerations*. WestEd.