

VAST Data Platform: Unified Authorization and Access Control with ABAC

Table of Contents

Introduction	3
Cloud IAM and ABAC	4
ZTA & MCS	6
VAST Data Platform	7
VAST Implementation of ABAC	8
Zero Trust Data Pillar	11
VAST Catalog	12
Audit	14
Conclusion	15

Introduction

Access and authorization vulnerabilities remain the Achilles' heel of cybersecurity, often serving as the root cause of most data breaches and exploits. Inadequate access controls can leave sensitive data exposed to unauthorized users, leading to significant security incidents. As cyber threats evolve, the need for robust access control mechanisms becomes increasingly critical. Traditional access control models, such as Role-Based Access Control (RBAC), have shown limitations in their ability to adapt to the complex and dynamic nature of modern IT environments, prompting the need for more advanced solutions.

VAST makes it possible to realize the benefits of Zero Trust Architecture (ZTA), Multi-Category Security (MCS), and Attribute-Based Access Control (ABAC) as a foundation for dynamic, modern data security practices, across cloud, data center and edge.

The next generation of access control began to take shape with the advent of more sophisticated database security systems. These systems aimed to address the shortcomings of RBAC by introducing fine-grained access controls that adapt to the changing context of user requests. It was from this evolution that Attribute-Based Access Control (ABAC) emerged as a promising solution. ABAC represents a paradigm shift in access control philosophy, moving beyond the static role assignments of RBAC to a more dynamic and context-aware model.

ABAC is a logical access control model that determines a user's authorization to perform operations by evaluating:

- Attributes associated with the Subject (requester or user)
- The Object (resource to be accessed)
- Requested operations such as Read, Modify, Delete
- In some cases, environmental conditions, such as location or time

This evaluation is conducted against policies, rules, or relationships that describe the allowable operations for a given set of attributes. For example, a user might have attributes such as role, department, and security clearance, while a file might have attributes like classification level, owner, and creation date. An ABAC policy could then allow access only if the user's clearance level is higher than the file's classification level and if the access request is made during working hours.

The flexibility of ABAC allows for the creation of highly granular and precise access control policies that can be evaluated and applied in real-time to the context of access requests. This level of detail and adaptability makes ABAC particularly well-suited for preventing unauthorized data access and mitigating the risk of data exploits. As organizations continue to grapple with the challenges of securing their data, ABAC stands out as a next-generation approach that provides the robust and flexible access control needed in today's complex security landscape.

Cloud IAM and ABAC

The advent of cloud computing has brought about a paradigm shift in identity, authorization, and access control, propelling these elements far beyond the capabilities of traditional on-premises deployments. The agility and innovation in cloud environments are largely due to the sophisticated approaches to managing digital identities and access permissions. Cloud providers like AWS, Azure, and Google Cloud have set the bar high with their advanced Identity and Access Management IAM systems, which offer fine-grained access control and support for ABAC, allowing permissions to be dynamically adjusted based on a multitude of attributes associated with users and resources.

AWS IAM, for instance, provides a robust framework for defining complex access rules in the widely adopted JSON format, which can include conditions based on both resource and user attributes. This flexibility is further enhanced by AWS's support for ABAC, where tags can be used as attributes to control access, enabling permissions to be fine-tuned in response to the attributes of the request or the resources involved.

Azure's IAM capabilities also support ABAC by leveraging Azure Active Directory for identity management. Azure allows for the creation of granular access policies in JSON format, which can be managed across multiple resources and subscriptions. This comprehensive approach to access control is essential for large enterprises that require sophisticated access management to secure their vast array of cloud resources.

Azure Cloud has further extended its IAM capabilities with Azure ABAC services through Azure Active Directory and Azure Storage. Azure ABAC builds upon RBAC by adding conditions to role assignments, using attributes associated with security principals, resources, requests, and the environment. This level of granularity in access control is crucial for maintaining data security and meeting compliance requirements.

Google Cloud IAM stands out with its robust access control framework, which supports ABAC through conditions in IAM policies. These conditions can be based on various attributes of the request, such as the identity of the requester, the time of the request, and the specific resource being accessed. Google Cloud's IAM policies, which can be defined in YAML or JSON format, enable precise control over resource access, ensuring that only authorized users can access the right resources under the right conditions.

Even in the realm of Linux distributions, such as Redhat, security modules like SELinux provide mechanisms for supporting access control security policies, including Mandatory Access Control(MAC). SELinux enhances the security of Linux systems by enforcing the separation of information based on confidentiality and integrity levels, addressing threats of tampering and attempts to bypass application security mechanisms.

While there are many commonalities of ABAC and IAM across the aforementioned clouds and Linux distributions, a lack of standardization means that each IAM system has its own specific syntax and services, which must be addressed in their respective IAM policies.

The primary reason for the emphasis on IAM policy management in the cloud is to enhance data security. Effective IAM policies are instrumental in protecting data privacy by limiting user access to resources, thereby preventing identity theft, data breaches, and unauthorized access to sensitive corporate information. IAM systems are essential for controlling access within a cloud environment, playing a critical role in ensuring that only authorized users can access specific resources. This is especially important in cloud-native environments where traditional network-layer controls may be insufficient, and services that don't expose underlying infrastructure rely heavily on IAM for security.

Moreover, the cloud's requirement for multiple tenancies, necessary for meeting regulatory requirements, means that each tenancy has its own set of IAM resources and rules, providing a high degree of operational autonomy. IAM policies are central to meeting regulatory compliance requirements and can help organizations adhere to industry regulations such as FedRAMP Impact levels 1 – 6, HIPAA, SOX, and PCI-DSS, which mandate the security and privacy of customer data.

ZTA & MCS

To address the complexities of on-premises deployments, the National Institute of Standards and Technology (NIST) is advocating for a more integrated approach through the adoption of Zero Trust Architecture (ZTA), Multi-Category Security (MCS), and Attribute-Based Access Control (ABAC). This initiative aims to modernize and secure on-prem environments with cloud-like frameworks and access control mechanisms. Zero Trust, a security model that operates on the principle of “never trust, always verify,” ensures that access to resources is granted based on continuous verification of identities and context, making it significantly harder for unauthorized users to gain access to sensitive data or systems.

ABAC, a critical component of ZTA, allows for a dynamic and granular approach to access control, where decisions are made based on a comprehensive evaluation of attributes related to the user, the resource, the environment, and the context of the access request. This allows for a more nuanced and flexible access control mechanism that can adapt to the complex security needs of large enterprises, especially those with disparate storage solutions. For instance, ABAC can enforce policies that restrict access to specific data sets or systems based on the user's role, location, device security posture, or even the sensitivity of the data being accessed, thereby enhancing and simplifying the security posture of on-prem deployments.

The integration of MCS within this framework ensures that resources are classified and protected based on their sensitivity and the potential impact of their compromise. By applying MCS principles, organizations can segment and protect data and resources across different categories, ensuring that access controls are appropriately stringent for higher sensitivity resources while allowing more flexible access to less critical systems. This segmentation is crucial for managing the diverse and often complex access requirements within large enterprises, enabling them to safeguard critical data against unauthorized access or breaches effectively.

However, the implementation of ZTA, ABAC, and MCS in on-prem environments presents significant challenges, particularly in terms of integrating these advanced access control mechanisms with legacy systems, which often rely on outdated and less flexible access control models like RBAC. The transition requires a thoughtful approach to system design, policy formulation, and technology integration, ensuring that new access control paradigms can work seamlessly with existing infrastructure.

Despite the potential of ZTA and MCS to bring cloud-like control and security to on-prem data centers, the greatest obstacle lies in the storage solutions. Legacy storage vendors typically do not offer the necessary feature sets to support these advanced frameworks, relying instead on disparate models such as RBAC for NAS and vendor specific IAM for object storage. This inconsistency across multi-vendor environments complicates the integration and uniform application of ZTA and MCS principles, making it challenging to achieve the desired level of control and security. The lack of compatibility with ZTA and MCS in legacy storage solutions not only hinders the transition to more modern access control models but also poses a significant risk to data security in an era where threats are increasingly sophisticated and pervasive.

VAST Data Platform

The VAST Data Platform stands at the forefront of data-intensive computing, heralding a new era as the AI operating system for data. It is a comprehensive software infrastructure meticulously crafted to capture, catalog, refine, enrich, and preserve data through real-time deep data analysis and deep learning. This innovative platform is engineered to provide seamless and universal data access and computing from edge-to-cloud, offering enterprises and cloud service providers the flexibility to deploy on their chosen infrastructure.

At the heart of the VAST Data Platform is the VAST DataStore, previously known as VAST' Universal Storage. The DataStore is a testament to the platform's ability to break the long-standing tradeoff between performance and capacity. It enables customers to manage their files, objects, and tables on a single tier of affordable flash storage, making it ready for AI scale data computing. The DataStore's ultimate purpose is to capture and serve the immense amount of natural raw, unstructured, and stream data that is being captured from the natural world.

The VAST DataBase, the platform's structured data management service, writes tables into the system and enables real-time, fine-grained queries into vast reserves of tabular data and cataloged metadata. It transcends conventional database management systems by combining the transactional capabilities of a row-based OLTP database with the analytical prowess of a flash-based data warehouse, all while maintaining the scale and affordability of a data lake.

The VAST Data Platform's architecture, known as DASE (Disaggregated, Shared-Everything), is a revolutionary distributed systems architecture that eliminates communication or interdependencies between the machines that run the system's logic. This architecture allows for scalable parallel read and write operations at any scale, with deep write buffers that enable data to be written in any form before being converted into their long-term data format on low-cost, high density flash. The platform can transact data in streams and then read and query data in ways optimized for deep data analytics and deep learning, breaking down the barrier between event-driven architectures and data-driven processing.

Furthermore, the VAST Data Platform is globally consistent and locally performant, introducing a new decentralized form of global transaction management that extends deep buffers and data synchronization across every location where data is computed and stored. Applications experience a strictly-consistent global namespace and dynamic function execution environment that extends from edge to cloud to on-premises data centers.

The VAST DASE architecture is designed to deliver 99.99999999% uptime in data centers on systems that can scale into the exabyte range. The platform is implemented with data access interfaces that are standard across enterprise applications, allowing VAST systems to integrate easily into environments. This modernizes the infrastructure of legacy applications, enables consolidation without application disruption, and brings AI to enterprise data, as opposed to shipping data to exotic AI data infrastructure.

VAST Implementation of ABAC




VAST's next-generation attribute-based access control (ABAC) solution authenticates protocols including NFSv4.1 with Kerberos, SMB 2.1 & 3, S3, and S3A. It seamlessly integrates with Microsoft Active Directory, representing a significant leap forward in data access control by leveraging user class attributes created as key-value pairs. This innovative approach, where a user-defined attribute name or data label like "apple=rw" denotes read-write access, revolutionizes access control by allowing complete management from the identity provider. This integration enables dynamic changes in access within the VAST platform without requiring direct access to the platform, meeting stringent separation of duty requirements because the attribute acts as an additional access control layer. As users request access to data, the system leverages existing role-based access control (RBAC) and protocol access control lists (ACLs). Once access is validated, a final check is performed against the user class attribute that matches the label(s) assigned to the object and/or file being accessed.

This capability works in conjunction with traditional access control mechanisms by considering not only the RBAC but also the specific tenant making the request, the network segment they are connected to, and the automated data labels (attributes) assigned to a view or path as data is created within the platform. Customers can define user class attributes within the identity provider, setting policies for read/write, read-only, and deny access. When a View (VAST's term for a multi-protocol capable data export) is created leveraging these attributes and a user accesses data, their identity, RBAC controls, and ABAC controls are evaluated to grant the appropriate level of access. VAST supports up to 20 different attributes for a single object and/or file stored within an ABAC-enabled View, ensuring that attributes are inherited for all files, objects, and directories within a view or encrypted path.

This hybrid approach between RBAC and ABAC, designed to meet NIST 800-53r5 Access Controls, assigns unique attributes as extended attributes (xattrs) for views, directories within views, and the objects/files within those views. When access is requested, the platform evaluates whether the attributes assigned to the View and object align with the defined access policy for that attribute.

Authentication & Authorization



To illustrate VAST's advanced ABAC solution, consider a scenario involving three users: cmccutcheon, rhays, and jdean, and user class attributes APPLE, GRAPE, and ORANGE. In this scenario, cmccutcheon is not authorized to access rhays's files, while jdean has access to both cmccutcheon's and rhays's files via traditional RBAC and ACL capabilities.

If jdean grants cmccutcheon access to rhays's data using traditional RBAC and ACL mechanisms, the VAST platform intervenes and denies cmccutcheon's access request when it is made. This is because rhays's files have an additional layer of security—an assigned attribute (APPLE)—and cmccutcheon lacks the corresponding read-only (RO) or read-write (RW) access defined within Active Directory for the APPLE user class attribute. This safeguard prevents unauthorized users, such as jdean, from improperly granting access to data, showcasing the VAST solution's superiority over legacy storage capabilities. Additionally, as this activity is captured and forwarded to the Network Operations Center (NOC), a Security Analyst can revoke jdean's access in Active Directory by simply changing the policy to DENY for all his attributes. The VAST Data Platform will then implement this change and deny jdean access to anything on the system.

The authentication and authorization process combines RBAC capabilities and layers on ABAC policies to fine-tune access permissions while enabling cloud-like control of access. This approach allows the simplicity of RBAC to be combined with the flexibility and granularity of ABAC, creating a robust and adaptable access control system that meets the evolving needs and security requirements of organizations, independent of the data protocol used for access. This is a key capability to simplify data pipelines where data must be accessed via multiple protocols, such as file and object, without compromising data security.

By integrating RBAC and ABAC, VAST's solution offers a powerful and flexible access control system that not only simplifies administration but also provides the granularity and adaptability required for changing organizational needs and security requirements. VAST's innovative approach to ABAC, with its tight integration with identity providers and dynamic access control capabilities, sets a new standard for data access control, highlighting the limitations of legacy storage solutions in meeting the demands of modern data-intensive applications.

VAST's ABAC solution also enables organizations to modernize access control for legacy applications without requiring extensive rewrites. For example, a content management system that uses SMB to store data can leverage VAST's ABAC capabilities to introduce next-generation authentication and authorization. By front-ending the legacy application with VAST's ABAC-enabled Views, organizations can provide secure access using ABAC while still allowing the application to use SMB with RBAC and ACLs for data storage. This approach extends the life and security of legacy applications, enabling a gradual modernization path without disrupting existing workflows. As a result, organizations can enhance the security posture of their legacy applications, meet evolving compliance requirements, and provide a more seamless and secure user experience by leveraging VAST's advanced ABAC solution.

Zero Trust Data Pillar

The VAST Data Platform's hybrid RBAC-ABAC model is a critical component in fulfilling the data pillar requirements for Zero Trust Architecture (ZTA). ZTA's core principle of "never trust, always verify" is embodied in the platform's ability to dynamically validate and authorize user access based on a comprehensive set of attributes. This approach ensures that trust is not assumed based on network location, but rather continuously earned through real-time verification against the organization's identity provider, such as Microsoft Active Directory. By incorporating ABAC, the VAST Data Platform can enforce granular access controls that align with the specific data security categories defined in MCS, allowing organizations to manage access to resources across different sensitivity levels and classifications. This is essential for maintaining a robust security posture and ensuring that data is accessed only by authorized personnel under the right conditions.

Bringing cloud-like control for data to any VAST deployment, the platform's ABAC solution enhances the traditional RBAC model by introducing automated data labeling. Every file or object within the system is automatically labeled with attributes that dictate access permissions, streamlining the enforcement of access policies. This automated labeling is particularly beneficial for legacy applications that traditionally operate on RBAC, ACL, and ACE models.

With VAST's ABAC layer, these applications can now benefit from an additional layer of security without any changes to their existing access control mechanisms. The ABAC model operates transparently, providing a seamless experience for users and applications while offering a more sophisticated and secure access control framework.

Additionally, the VAST Data Platform ensures that users cannot view the attributes associated with data, overcoming known exploits in MCS and MLS systems that arise from such visibility. By concealing these attributes, VAST not only enhances security but also maintains the integrity of the separation of duties, ensuring that even system administrators with high-level access cannot override the ABAC-enforced policies.

VAST Catalog

The VAST Catalog, a key feature of the VAST Data Platform, revolutionizes metadata management by presenting the metadata of the VAST Element Store file/object namespace as a database table. This innovative approach eliminates the need for resource-intensive file system tree walks traditionally required for metadata searches. Instead, system operators can efficiently query the VAST Catalog, which is populated from a VAST snapshot to provide a consistent point in time view. The catalog is updated as frequently as every 15 seconds, enabling point-in-time queries that reflect not just the current state of the data but also any indexed historical point. Queries to the VAST Catalog can be made through the VAST UI or via SQL-compatible tools, facilitating the discovery of data spillage events and enhancing data governance.



Index System and User Metadata

Metadata index built on native tabular elements



Faster, simpler search

vs POSIX or Object commands without file system crawl



Exascale Data Management

Scalable, self-managing, always in sync

System Metadata

- Time (ctime, mtime, atime)
- File name, object key
- File / Object size
- Ownership (User / Group)
- Path / Directory, bucket
- ABAC XATTR

User Metadata

- S3 Tags

Point in Time Views

- Built on VAST Snapshots
- Query any point in time



VAST Catalog



VAST Catalog metadata index is fully integrated into the VAST Data Platform

The VAST Catalog's metadata indexing is fully integrated into the VAST Data Platform, supporting NIST controls in several critical areas:



Access Control (AC)

The cataloging and indexing of data assets aid in access control by providing a comprehensive inventory of data that needs protection. This helps in determining appropriate access levels for different users and ensuring that access controls are effectively implemented.



Audit and Accountability (AU)

By maintaining a detailed record of data access, including who accessed what data and when, the VAST Catalog supports audit and accountability. This level of transparency is essential for identifying potential security incidents and verifying appropriate data usage.



Configuration Management (CM)

The VAST Catalog assists in configuration management by offering a clear inventory of data assets and their attributes. This ensures that data management aligns with security policies and procedures.



Identification and Authentication (IA)

The VAST Catalog enhances identification and authentication processes by providing a quick and reliable method to locate and verify the authenticity of data, helping to prevent unauthorized alterations.

The VAST Catalog's ability to index, search, and catalog data plays a pivotal role in meeting NIST controls, offering a clear view of data assets and their attributes, which is crucial for managing access, audits, configuration, identification, and authentication more effectively. In comparison to cloud IAM and ABAC capabilities, the VAST Data Platform offers a unique advantage. It allows customers to utilize VAST to bring legacy datasets to the cloud without undergoing a full cloud migration, which typically involves developing new access control strategies. On-premises data and applications that leverage RBAC with ACLs can now be managed with cloud-like IAM policies and ABAC capabilities, thanks to VAST's platform. This means that files traditionally treated as objects in the cloud can be seamlessly managed on-premises with the same level of control and security, bridging the gap between legacy and modern cloud environments.

Audit

The VAST Data Platform's auditing capabilities are designed to provide a comprehensive and compliant framework that aligns with the NIST SP 800-53 Revision 5 Audit and Accountability (AU) controls. The platform employs a trio of auditing frameworks—Auditd (Audit Daemon), Admin, and Protocol Audit—to capture and log a wide range of audit events, ensuring that all system, protocol, and administrative activities are meticulously recorded.

The Auditd framework with DISA STIG rules is responsible for monitoring and capturing system-level events, providing a detailed log of actions taken within the platform's operating environment. This includes tracking user logins, file access, and system changes, which are crucial for maintaining a secure and traceable record of operations.

The Admin audit framework focuses on administrative activities, offering insights into the actions performed by system administrators. This includes changes to user permissions, system configurations, and policy updates. By logging these activities, the platform ensures that any administrative action is accountable and can be reviewed for compliance and security purposes.

The Protocol Audit framework captures events related to the various communication protocols used within the platform. This includes data transfers, access requests, and protocol-specific interactions that occur during the operation of the platform. By auditing these events, the VAST Data Platform can ensure the integrity and security of data in transit and provide a clear trail of protocol-based activities.




VAST's robust audit capabilities are fully compatible with popular SIEM platforms

Together, these frameworks work in concert to capture a holistic view of audit events across the platform. As events are logged, they can be forwarded to a Security Information and Event Management (SIEM) system, allowing organizations to analyze and respond to potential security incidents effectively.

The VAST Data Platform satisfies several NIST SP 800-53 Revision 5 Audit and Accountability (AU) 2 - 12 controls with its auditing functionality that ensures organizations can effectively manage audit generation, storage, analysis, and review, providing a robust framework for audit and accountability.

Conclusion



In conclusion, the VAST Data Platform stands as a pioneering force in the realm of data management and security, seamlessly blending the robustness of Role-Based Access Control (RBAC) with the dynamic flexibility of Attribute-Based Access Control (ABAC). This integration not only adheres to the NIST Special Publication (SP) 800-162 guidelines for ABAC but also propels the platform to the forefront of addressing modern cybersecurity challenges. By leveraging user class attributes created as key-value pairs denote access permissions, VAST revolutionizes access control. This approach allows for complete management from the identity provider, enabling dynamic changes in access within the VAST platform and ensuring compliance with stringent separation of duty requirements.

The platform's auditing capabilities, utilizing Auditd, Admin, and Protocol Audit frameworks, capture and log a wide range of audit events, aligning with the NIST SP 800-53 Revision 5 Audit and Accountability (AU) controls. This comprehensive auditing framework ensures meticulous recording of all system, protocol, and administrative activities, enhancing an organization's security posture and response capabilities. Combining Vast Catalog and Audit information enables customers to in real-time detect who did what, where, and when on the platform.

Moreover, the VAST Catalog transforms metadata management by presenting the metadata of the VAST Element Store file/object namespace as a database table, enabling efficient queries without resource-intensive tree walks. Now you can search, categories, and visualize data within the platform independent of how you access that data. These features, coupled with the platform's automated data labeling capabilities, brings cloud-like control for data to any deployment. It enhances traditional RBAC models with its ABAC solution, providing a seamless experience for users and applications while offering a sophisticated and secure access control framework.

The VAST Data Platform not only meets the rigorous demands of Zero Trust Architecture (ZTA) and Multi-Category Security (MCS) but also bridges the gap between on-premises legacy NAS data and cloud IAM capabilities. Organizations can leverage VAST to modernize their data access control strategies without undergoing extensive reconfiguration or migration, highlighting the platform's innovative approach to ABAC, its tight integration with identity providers, and its dynamic access control capabilities.

In essence, the VAST Data Platform sets a new standard for data access control, highlighting the limitations of legacy storage solutions in meeting the demands of modern data-intensive applications. Its hybrid RBAC-ABAC model is a pivotal advancement in access control, providing automated data labeling and attribute-based access control that brings sophisticated, cloud-like data management capabilities to any deployment. Reinforcing the platform's position as a leader in the next generation of data security solutions, the VAST Data Platform is poised to transform the landscape of data access control, ensuring that organizations can navigate today's complex security landscape with confidence and agility.



For more information on the VAST Data Platform and how it can help you solve your application problems, reach out to us at hello@vastdata.com.