



QUBOLE ON MICROSOFT AZURE

SECURITY AND COMPLIANCE
WHITE PAPER



© 2018 Qubole. All rights reserved.

Qubole

469 El Camino Real, Suite 201

Santa Clara, CA 95050

www.qubole.com

(855) 423-6674

Trademarks

Qubole is a registered trademark. All other trademarks or registered trademarks are the property of their respective owners.

Disclaimer

The information provided in this document is provided “as is” without warranty of any kind. Qubole disclaims all warranties, either express or implied, including the warranties of merchantability and fitness for a particular purpose. In no event shall Qubole be liable for any damages whatsoever including direct, indirect, incidental, consequential, loss of business profits or special damages, even if Qubole or its suppliers have been advised of the possibility of such damages.

Document Lifetime

Qubole may occasionally update online documentation between releases of the related software. Consequently, if this document was not downloaded recently, it may not contain the most up-to-date information. Please refer to www.qubole.com for the most current information.

Where to get help

Qubole support, product, and licensing information can be obtained as follows.

Product information — For documentation, release notes, software updates, or for information about Qubole products, licensing, and service, see the Qubole website at:

<http://www.qubole.com>

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send your opinion of this document to: support@Qubole.com

If you have issues, comments, or questions about specific information or procedures, please include the title and, if available, the part number, the revision, the page numbers, and any other details that will help us locate the subject that you are addressing.

TABLE OF CONTENTS

INTRODUCTION	5
ORGANIZATIONAL SECURITY—OUR PHILOSOPHY	6
PERSONNEL SECURITY	6
SECURITY AND PRIVACY TRAINING	6
Product Security	6
Security Operations	6
CSIRT	6
Risk and Compliance	6
QUBOLE'S SECURITY PROFESSIONALS	7
POLICIES AND STANDARDS	8
AUDITS, COMPLIANCE, AND THIRD-PARTY ASSESSMENTS	9
Cloud Security Alliance	9
AICPA Service Organization Controls 2 Type 2 (SOC2)	9
Privacy Shield	9
GDPR	10
ISO/IEC 27001	10
Penetration Testing	11
DESIGN SECURITY	11
PROTECTING CUSTOMER DATA	12
DATA ENCRYPTION	12
Data in Transit	12
Data at Rest	12
PLATFORM SECURITY	13
Firewalling via Network Security Groups	13

Security Event and Incident Management	13
Intrusion Detection	13
File Integrity	13
Vulnerability Scanning and Detection	14
Hardened Baselines and Configuration Standards	14
AUTHENTICATION	14
MONITORING SYSTEMS, SERVERS, AND SETTINGS	15
System Monitoring, Logging, and Alerts	15
Continuous Development and Improvement	15
Controlling Change	15
Preventing and Detecting Malicious Code	15
SERVER HARDENING	16
DISASTER RECOVERY AND BUSINESS CONTINUITY	16
THIRD-PARTY SUPPLIERS	16
QUBOLE ON MICROSOFT AZURE	17
Hive Authorization in Azure —Improving Enterprise-level Security and Data Governance in the Cloud	17
Creating an Account and Authorizing Users	17
Single Sign-On	17
OAuth	18
SAML	18
Active Directory Federated Service (ADFS)	18
ACCESSING DATA SECURELY	19
How Azure RBAC Roles Work to Manage Secure Access and Authorization	20
Per-User API Tokens	20
LEARN MORE	21
QUBOLE ON MICROSOFT AZURE	21



INTRODUCTION

Qubole provides a cloud-native data platform for self-service artificial intelligence (AI), machine learning (ML), and analytics, and is committed to providing easily accessible data-driven insights. Our customers process nearly an exabyte of data every month—on Microsoft Azure, Amazon Web Services, or Oracle Cloud—so we know that safeguarding that data, enforcing proper security measures, ensuring regulatory compliance, and providing trusted services are essential to your success. This paper discusses the security strategies we use to protect your information and provides details of how that strategy is implemented on Microsoft Azure.





ORGANIZATIONAL SECURITY— OUR PHILOSOPHY

We believe that availability, security, governance, and privacy are essential elements of the Qubole Data Service (QDS). Security is integral to our culture and our ongoing mission. Security awareness at Qubole far exceeds the minimal standards required by auditors and compliance organizations and is incorporated in all aspects of our business model. From product design and development to day-to-day corporate operations, we are constantly building and improving our robust organizational security programs.



PERSONNEL SECURITY

Qubole's personnel practices apply to all members of our workforce —regular employees and contractors—who have access to information systems. All Qubolers are required to understand and follow internal policies and standards. Before gaining initial access to systems, Qubolers must agree to confidentiality terms, pass a rigorous background screening, and thereafter, must complete additional security training annually. Upon termination, all access to our systems is immediately removed.



SECURITY AND PRIVACY TRAINING

During their tenure, all members of our workforce are required to complete a privacy and security training refresher course at least once per year. These courses cover privacy and security topics, including device security, acceptable use, preventing malware, physical security, data privacy, account management, and incident reporting. Qubolers are also required, on an annual basis, to acknowledge that they've read and will follow Qubole's information security policies.

Certain professionals, such as engineers, operators, and support personnel who may have elevated access to systems or data, receive additional job-specific training on privacy and security. Qubolers are also required to report security and privacy issues to appropriate security teams.



QUBOLE'S SECURITY PROFESSIONALS

Qubole employs a Chief Information Security Officer (CISO) who functions as the organization's top executive responsible for security, compliance, and privacy. Our CISO serves as the company's business leader responsible for developing, implementing, and managing the organization's corporate security vision, strategy, and programs.

The CISO is supported by a team of security professionals with more than 75-years of combined enterprise security and privacy experience. This team focuses on Product security, security operations, computer security incident response, risk, and compliance. Together, they share responsibilities for key aspects of our security program including:

Product Security

- Establishing secure development practices and standards
- Conducting security risk assessments
- Providing design and code review recommendations for detecting and removing common security flaws
- Training developers on secure coding practices, data privacy, and governance and risk mitigation strategies

Security Operations

- Building and operating critical security infrastructure including Qubole's public key infrastructure, event monitoring, and authentication services
- Maintaining a secure archive of security-relevant logs
- Consulting with operations personnel to ensure the secure configuration and maintenance of Qubole's production environments

Computer Security Incident Response Team (CSIRT)

- Responding to alerts related to security events on Qubole systems
- Managing security incidents
- Acquiring and analyzing threat intelligence

Risk and Compliance

- Coordinating penetration testing
- Managing vulnerability scanning and remediation
- Coordinating regular risk assessments, and defining and tracking risk treatments
- Managing the security awareness program
- Coordinating audits and maintaining security certifications
- Responding to customer inquiries
- Reviewing and qualifying vendors' security posture

All members of Qubole's Security Team are active participants in the larger information security community to improve the overall state of the art of information security and to maintain their own expertise.



POLICIES AND STANDARDS

Qubole maintains a set of policies, standards, procedures, and guidelines ("security documents") that provide our workforce with the "rules of the road." Our security documents help ensure that Qubole customers can rely on us to behave ethically and for our service to operate securely. Security documents include, but are not limited to:

- Fair, ethical, and legal standards of business conduct and acceptable use of systems
- Classification, labelling, and handling rules for all types of information assets
- Practices for worker identification, authentication, and authorization for accessing system data and secure development, acquisition, configuration, and maintenance of systems
- Workforce requirements for transitions, training, and compliance with Information Security Management System (ISMS) policies
- Use of encryption, including requirements for when and where to use it
- Description, schedule, and requirements for retention of security records and planning for business continuity and disaster recovery
- Classification and management of security incidents and change control
- Regular use of security assessments such as risk assessments, audits, and penetration tests



AUDITS, COMPLIANCE, AND THIRD-PARTY ASSESSMENTS

Qubole evaluates the design and operation of its overall ISMS for compliance with internal and external standards. We engage credentialed assessors to perform external audits at least once per year including the following organizations:

Cloud Security Alliance

Qubole has completed the Consensus Assessment Initiative Questionnaire (CAIQ). This program outlines more than 200 questions relating to cloud services, provisioning, storage, security, availability, confidentiality, and privacy. The current Cloud Security Alliance CAIQ is free to download and can be found [here](#)¹. To learn more about Qubole's responses to the CAIQ, or to obtain a copy of this document please reach out to your Qubole Sales representative.

AICPA Service Organization Controls 2 Type 2 (SOC2)

Qubole has successfully completed the SOC 2 Type II examination. The report generated from this extensive evaluation was conducted by an American Institute of CPAs (AICPA) accredited firm, which attests that Qubole meets stringent guidelines for data security.

SOC 2 reports are designed to provide information and assurances that the service delivery processes and controls used by a SaaS provider meet certain high standards and guidelines for data security, confidentiality, and availability. There are two types of SOC 2 reports: Type I and Type II. We opted for the much more rigorous Type II report because security is of the utmost importance to us and to our many clients who entrust us to power their AI, ML, and analytics initiatives in the cloud.

Note: Qubole requires that you sign a Non-Disclosure Agreement to receive a copy of this SOC2 Type II report. Please reach out to your Qubole Sales representative for more information.

Privacy Shield

Qubole has engaged with TrustArc (formerly TRUSTe) to complete its assessments and attest to Qubole's compliance with the US Privacy Shield regulation for privacy and transfer of European Union (EU) Personal Data to the United States. Qubole currently self-certifies with the Department of Commerce International Trade Administration (ITA).

¹ <https://cloudsecurityalliance.org/artifacts/consensus-assessments-initiative-questionnaire-v3-0-1/>

General Data Protection Regulation (GDPR)

GDPR is a European Data Privacy legislation that became effective in May 2018. This regulation outlines the importance of data subjects' rights to the privacy of their personally identifiable information (PII). Qubole fully supports and welcomes the GDPR as an opportunity to deepen its commitment to data protection. Qubole complies with the GDPR in the delivery of our service to our customers. Qubole will also continue to enhance data protection and compliance in the following areas:

1. Accountability, policies, and procedures
2. Compliance and risk activities
3. Implementing additional security measures
4. Notification and reporting requirements for data breaches

Qubole has partnered with a Third Parties to continue to align innovation in big data for AI, ML, and analytics to meet the exacting requirements of the GDPR. In addition, Qubole has created a Data Processing Addendum as an attachment to its Master Services Agreement. This document supports our commitment to this important legislation and is available for download at <https://trust.qubole.com/#legalcompliance>

ISO/IEC 27001

The ISO/IEC 27000 family of standards helps organizations keep information assets secure. ISO/IEC 27001 is the best-known standard in the family providing requirements for an information security management system (ISMS). ISO/IEC 27001 specifies the requirements for establishing, implementing, maintaining and continually improving an ISMS within the context of the organization. It also includes requirements for the assessment and treatment of information security risks tailored to the needs of the organization.

Qubole believes that compliance must be built throughout QDS and, as such, our processes, procedures, controls, operations, and activities align with the ISO/IEC 27001 standards and are reflected in our policies and other attestation and certification work. Qubole is currently certifying for ISO/IEC 27001, so please reach out to your Qubole Sales representative for more information.

Audit results are shared with senior management and all findings are tracked to resolution

Penetration Testing

Qubole engages independent entities to conduct regular application and infrastructure-level penetration tests. Results of these tests are shared with senior management. Qubole's Security Team reviews and prioritizes the reported findings create tasks for engineering to resolve any deficiencies and tracks them to resolution.

Qubole will share a summary of the most recently performed tests under non-disclosure upon request; please reach out to your Qubole account representative. Customers wishing to conduct their own penetration test of Qubole's system can contact their account representative regarding authorization to perform this test.

Qubole aligns our penetration test against OWASP Top Ten. In summary, some of the areas covered are listed below:

1. Injection Flaws – Qubole not only evaluates these flaws during the penetration tests but also uses Static Code Analysis Testing during development.
2. Broken Authentication – Qubole uses mechanisms to strengthen and securely store passwords, secrets, and tokens.
3. Security Misconfiguration – Qubole removes default components including default credentials. Qubole turns off and disables default services and hardens the baseline of compute systems.
4. Cross Site Scripting (XSS) – Qubole evaluates the three primary types of XSS and has libraries in place (in code) to sanitize and prevent exploitation.
5. Insufficient Logging and Monitoring – Qubole centrally logs all critical system and application logs. Qubole also has built-in anomaly detection for questionable events and actions and monitors for abuse and attacks through log analytics.



DESIGN SECURITY

Qubole assesses the security risk of each software development project according to our Secure Software Development Lifecycle. Before completion of the design phase, we undertake an assessment to qualify the security risk of the software changes introduced. This risk analysis leverages both the OWASP Top 10 and the experience of Qubole's Product Security team to categorize every project as High, Medium, or Low risk.

Based on this analysis, Qubole creates a set of requirements that must be met before the resulting change may be released to production. All code is checked into a version-controlled repository. Code changes are subject to peer review and continuous integration testing. Qubole's Security Team also operates continuous automated static analysis using advanced tools and techniques. Significant defects identified by this process are reviewed and followed to resolution by the Security Team.



PROTECTING CUSTOMER DATA

The focus of Qubole's security program is to prevent unauthorized access to customer data. To this end, our team of dedicated security practitioners, working in partnership with peers across all our teams, take exhaustive steps to identify and mitigate risks, implement best practices, and constantly evaluate ways to improve our security measures.



DATA ENCRYPTION

Qubole supports end-to-end encryption throughout Qubole services as defined below. Encryption protects the confidentiality of customer data ensuring that only those users and entities with proper permissions can access data (Note that encryption could impact QDS performance).

Data in Transit

Qubole transmits data that travels over public networks using strong encryption. This includes data transmitted between Qubole clients and the service. Qubole also supports Transport Layer Security (TLS) encryption in the communication between members of the clusters, between the various service offerings, and between clusters and the Azure Blob Storage.

Qubole supports the latest recommended secure cipher suites to encrypt all traffic in transit, including use of TLS v1.2 for communication between customer browsers and clients, REST API endpoints, and Qubole services. Qubole has enabled this by default.

Qubole monitors the changing cryptographic landscape and upgrades the cipher suite choices as the landscape changes, while also balancing the need for compatibility with older browsers and endpoints.

Data at Rest

Qubole supports server-side encryption on Azure Blob Storage and Azure Data Lake Store (ADLS). Encryption is enabled by default in both the Blob Storage and ADLS with keys managed by the service. Customers also have the option to provide their own encryption keys managed in Azure Key Vault. Qubole works seamlessly with both these options without needing any user-side configuration.

PLATFORM SECURITY

Firewalling via Network Security Group

Qubole applies strict security measures when creating and maintaining customer clusters. Qubole dynamically creates the rules that are necessary to access and support your clusters. We also create encryption access keys to clusters and dynamically encrypt ephemeral storage with one-time-use keys that are destroyed on cluster terminate, and are never stored. This is enforced—at a minimum—with secure permissions, granting access from the Qubole Orchestration Tier to the Compute/Data Plane, which resides in the customer's Microsoft Azure account. All communication between the Orchestration Plane and the Data Plane is encrypted. Qubole also supports the use of Private IPs for clusters. In this case, all communication between the Qubole Orchestration Tier and the Compute/Data Plane is supported via a Bastion node, so customers can use Private IPs for all cluster nodes with a public IP required only for the Bastion node.

Security Event and Incident Management

Qubole maintains a Security Event and Incident Management System (SEIM) which is a central platform where alerts from a variety of sources are forwarded for review. In the event an event exceeds a given threshold, the event is forwarded to an on-call engineer for evaluation and escalation. This monitoring learns patterns which over time can predict nominal behavior and identify behaviors that may be suspicious and look for and alert on that behavior.

Intrusion Detection

Qubole deploys a log and kernel-based intrusion detection system that captures system changes, anomalous behavior and identifies anomalies and spurious behavior on instances within the Qubole Orchestration Tier. When there is an unusual event detected in system or audit logs or in commands issued or series unusual events, alarms are triggered in accordance with industry best practices and are forwarded to the SEIM or to an on-call engineer.

Qubole supports customer supplied and operated Intrusion Detection (IDS) or Intrusion Prevention (IPS) services within the customer account via bootstrapping and usage of custom images. Your solutions architect can explain how this works and you can find additional documentation here: [Understanding a Node Bootstrap Script](#)

File Integrity

Qubole operates a file integrity (FIM) management service that detects changes and acts as a defence against system compromise, malicious behaviour, and unauthorized actions by monitoring certain critical files on instances within the Qubole Orchestration Tier such as changes to critical logs, Qubole application code, SetUID binaries and additional configuration files and objects that may indicate a host has had a configuration change that should be investigated. When an unusual file event is detected, an alert is forwarded to the Qubole SEIM.

Vulnerability Scanning and Detection

Qubole externally and internally scans systems to determine if there are any application or system-specific risks. Qubole has a dedicated operations team who work around-the-clock to remediate these risks. Vulnerability detection is used to identify known software vulnerabilities in the packages, applications, operating systems, databases, and services used in QDS. These scans are configured to run bi-weekly and any identified vulnerabilities based on industry best practices based on severity.

Hardened Baselines and Configuration Standards

Qubole evaluates the baseline of systems deployed in the Qubole Orchestration Tier using the Center for Internet Security (CIS) benchmark. This benchmark covers not only applications installed, configured, and running on systems, but also system and configuration files including restrictions around permissions, logs, and system binaries. This benchmark covers more than 150 discrete changes and is applied to all instances deployed to both the Qubole Orchestration Tier as well as instances launched in customer accounts via the Qubole Data Plane Cluster Image.



AUTHENTICATION

To further reduce the risk of unauthorized data access, Qubole employs multi-factor authentication for all administrative access to systems with more highly sensitive and regulated data. Where possible and appropriate, Qubole uses unique private keys for authentication. For example, at this time, administrative access to production servers requires operators to connect from a known Internet Protocol address, VPN and using both SSH and a one-time password associated with a device-specific token.

Where passwords are used, multi-factor authentication is enabled for access to sensitive and regulated data. Passwords are required to be complex. System passwords are auto-generated to ensure uniqueness. These system passwords and accounts must have passwords longer than 20 characters and cannot consist of a single dictionary word (among other requirements). Qubole requires personnel to use an approved password manager. Password managers generate, store and enter unique and complex passwords. Using a password manager helps avoid password reuse, phishing, and other behaviors that can reduce security.



MONITORING SYSTEMS, SERVERS, AND SETTINGS

Qubole monitors and analyses the complete operating environment, ensuring that no potential security issues remain unexamined or unresolved.

System Monitoring, Logging, and Alerts

Qubole monitors servers and workstations to retain and analyze a comprehensive view of the security state of its corporate and production infrastructure. Administrative access, use of privileged commands, and system calls on all servers in Qubole's production network are logged. Qubole's Security Team collects and stores production logs for analysis. Logs are stored in a separate network. Access to this network is restricted to members of the Security Team. Logs are protected from modification and retained for two years. Log analysis is automated to the extent practical to detect potential issues and alert responsible personnel. Alerts are examined and resolved based on documented priority.

Continuous Development and Improvement

We proactively take a variety of steps to combat the introduction of malicious or erroneous code to our operating environment and protect against unauthorized access. This is done through multiple steps including security architecture reviews, application, security and privacy impact assessments, peer reviews, automated and manual static code analysis and other techniques.

Controlling Change

To minimize the risk of data exposure, Qubole controls changes, especially changes to production systems, very carefully. Qubole applies change control requirements to systems that store data at higher levels of sensitivity. These requirements are designed to ensure that changes potentially impacting customer data are documented, tested, and approved before deployment.

Preventing and Detecting Malicious Code

In addition to general change control procedures that apply to our systems, Qubole production network is subject to additional safeguards against malware.



SERVER HARDENING

New servers deployed to production are hardened by disabling unneeded and potentially insecure services, removing default passwords, unnecessary system accounts, and applying Qubole's custom configuration settings to each server before use. Qubole performs hundreds of additional changes to the default image to ensure that security is paramount and the systems are protected from unauthorized access.



DISASTER RECOVERY AND BUSINESS CONTINUITY

Qubole uses services provided by its cloud provider to distribute its production operation across multiple physical locations. These locations are within one geographic region, but protect Qubole service from loss of connectivity, power infrastructure, and other common location-specific failures. Production transactions are replicated among these discrete operating environments, to protect the availability of Qubole service in the event of a location-specific catastrophic event. Qubole also retains a full backup copy of production data in a remote location more than 2500 miles from the location of the primary operating environment. Full backups are saved to this remote location once per day, and transactions are saved continuously. Qubole tests backups at least quarterly to ensure they can be correctly restored.



THIRD-PARTY SUPPLIERS

To run its business efficiently, Qubole relies on sub-service organizations. Where those sub-service organizations may impact the security of Qubole production environment, Qubole takes appropriate steps to ensure its security posture is maintained. Qubole establishes agreements that require service organizations to adhere to confidentiality commitments Qubole has made to its users. Qubole monitors the effective operation of the organization's safeguards by conducting reviews of its service organization controls before use and at least annually.

QUBOLE ON MICROSOFT AZURE

Natively designed for Microsoft Azure and tightly integrated with its Blob storage, ADLS, compute, and other key architectural elements, running QDS on Azure provides unparalleled enterprise-level security and data governance. Account provisioning is straightforward and, unlike other big data-as-a-service providers, Qubole does not charge you to set up enterprise authentication and authorization that supports all your business needs.

Defining Your Security Options in Azure

Setting up a Qubole account and leveraging Azure security features is a simple process and provides virtually unlimited flexibility. Qubole supports SAML, OAuth, ADFS, whitelisting and granular access controls and, while it is your responsibility to create and define your accounts, grant access to data, and configure your cloud services, Qubole support is always available to assist you.

Hive Authorization in Azure —Improving Enterprise-level Security and Data Governance in the Cloud

Qubole supports data-centric security with Hive authorization. This improves usability for both cloud-storage administrators and data administrators while eliminating errors that arise from end-user authorization problems and by providing granular data access controls. This is an important milestone that just adds to Qubole's commitment of building and maintaining a secure, enterprise-level cloud platform.

Creating an Account and Authenticating Users

To create a Qubole account, visit Qubole at https://azure.qubole.com/users/sign_up. You have the option to sign up using email or your google account. Sign-up allows you to test drive Qubole on Azure for free without incurring any cloud or Qubole charges. You can try out examples we have built for data analysts and data scientists in our example gallery using Hive, Spark or Presto engines. Test drive is limited to 14 days after which you have the option to upgrade to an enterprise edition or unlock a free 30-day trial where you can preview additional Qubole features using your own datasets.

Single Sign-On

Enterprise Ready



OAuth 2.0



SAML v2



ADFS

Qubole supports all industry-standard authentication and authorization protocols including OAuth 2.0, SAML 2.0 and ADFS. You can also use the Qubole user management features to administer your Qubole users. To locate this feature in Qubole, navigate to Control Panel > Manage Users.

OAuth

Open Authentication (OAuth) can be configured for use with Qubole. OAuth extends a token validated against another form of authorization like Google. Qubole supports Google OAuth 2.0, Microsoft OAuth 2.0. To use OAuth please find a reference here for more detailed setup instructions: <http://docs.qubole.com/en/latest/admin-guide/saml-sso.html>.

SAML

Security Assertion Markup Language (SAML) is an enterprise authentication and authorization schema that provides for centralized authentication and provisioning of users from a secondary source that can be controlled by the enterprise. SAML is an open, XML-based standard that can be configured using several third parties including Okta, OneLogin, LastPass, Ping Identity and others who support SAML v2.0.

Active Directory Federated Service (ADFS)

ADFS allows an enterprise to connect to a corporate Active Directory to enable authentication that aligns with existing corporate needs. Qubole fully supports single sign-on to the platform through integration with ADFS or Azure Active Directory via a SAML 2.0 token. To use SAML please find a reference here for more detailed setup instructions: <http://docs.qubole.com/en/latest/admin-guide/saml-sso.html>.



ACCESSING DATA SECURELY

Qubole takes advantage of Active Directory RBAC roles in Azure to limit access to resources such as storage and compute. Using a refined set of permissions allows our customers to use Qubole on their behalf. Qubole also provides the right permissions for your data analysts and operators to create, run and modify queries and commands. Qubole's extensive functionality allows you to configure fine-grained access controls and permissions across Qubole resources.

Additionally, common concerns are addressed including limiting access rights to modify or affect the status of clusters, limiting the types of commands your users can execute and the data engines they can use. For more information, see

[Managing Roles in QDS](#)

Manage Roles +				
Role Name ↕	Policy			Actions
	Access	Resource	Policy Actions	
dashboard-user	allow	Clusters	start	▼
	allow	Notebook Dashboards	read, execute	
	allow	Folder	read	
system-admin	allow	All	all	▼
system-user	allow	All	read	▼
	allow	Commands	create	
	allow	Clusters	start	
	allow	Templates	create, clone, run	
	allow	Workspace	create, update	
	allow	Account	auth_token	
	allow	Scheduler	create, clone	
	allow	Scheduler Instance	kill, rerun	
	allow	App	create	
	deny	Data Store	create, update, delete	

How Azure RBAC Roles Work to Manage Secure Access and Authorization

Qubole uses Azure RBAC roles to securely access compute and storage in the customers Azure tenant. Azure RBAC is an authorization system that provides fine-grained access management to Azure resources, such as compute and storage.

Qubole uses Azure RBAC roles to limit the privileges granted to Qubole within a customer's Azure tenant. Specifically, customers need to set up an Azure Active Directory application for Qubole and grant the application either a contributor role or create a custom role that provides sufficient permissions to the application to create and terminate clusters in the customer's Azure tenant.




You can grant Qubole clusters access to Azure blob storage using a storage account key. Access to Azure Data Lake Store (ADLS) can be done either by creating a separate Active Directory Service Principal for ADLS or by using the same one created for compute access.

Qubole will initiate and terminate instances on your behalf based on your configuration which you set inside of the QDS Control Panel. This includes minimum and maximum cluster size, and on-demand node types to form heterogeneous clusters in the region specified by the customer.

Per-User API Tokens

Each user can also be configured with their own API token allowing granular access controls and auditing on a per-user basis. This provides you with maximum control over your users when used in conjunction with the roles described above. This token can also be reset to comply with more stringent security controls.

You can locate your (user) authentication token by navigating in QDS to Control Panel > My Accounts > Show API Token.

 My Accounts				
Account Name ↕	Account Id ↕	Status ↕	My Groups ↕	API Token
Acme Widgets 	3030		system-admin, system-user	Show Reset



CONCLUSION

We take security seriously at Qubole and we pride ourselves on the vigilance we employ to protect our customers' data assets. We believe that a mature security organization requires coordinated dedication across technology, policy, procedures, and people. This dedication is underscored by the risk-based approach laid out in this document to demonstrate strength at every layer of security. From compliance audits to HIVE authorization in Azure, our security strategy is designed to minimize any potential vulnerability or weakness. We want our customers to know their data is sufficiently protected, and we welcome the opportunity to discuss our security practices further.



LEARN MORE

For the latest information about our product and services, please see the following resources:

- [Qubole Whitepapers](#)
- [Qubole Case Studies](#)
- [Qubole Technical Documentation](#)

Connect with us:

469 El Camino Real, Suite 201

Santa Clara, CA 95050

www.qubole.com

(855) 423-6674

About Qubole

Qubole is revolutionizing the way companies activate their data — the process of putting data into active use across their organizations. With Qubole's cloud-native big data platform, companies exponentially activate petabytes of data faster, for everyone and any use case, while continuously lowering costs. Qubole overcomes the challenges of expanding users, use cases, and variety and volume of data while constrained by limited budgets and a global shortage of big data skills. Qubole offers the only platform that delivers freedom of choice, eliminating legacy lock in — use any engine, any tool, and any cloud to match your company's needs. Qubole investors include CRV, Harmony Partners, IVP, Lightspeed Venture Partners, Norwest Venture Partners, and Singtel Innov8.

For more information visit www.qubole.com.