Introduction

Qubole believes that availability, security, governance and privacy are essential and integral elements of the design, development and operations of the Qubole Data Service® (QDS).

Qubole currently performs bi-quarterly penetration tests, statically analyzes application infrastructure and cloud environments for vulnerabilities and risk, and evaluates all parties involved in the service supply chain.

Our goal is to make Qubole your trusted big-data partner. You can find more information on this at our *trust* site located at https://trust.qubole.com.

Compliance

Cloud Security Alliance

Qubole has completed the Consensus Assessment Initiative which outlines more than 200 questions that relate to cloud services, provisioning, storage, security, availability, confidentiality and privacy. A copy of this report can be requested from your account manager or account executive. The current CSA CAIQ report can be found here and is free to download.

AICPA Service Organization Controls 2 Type 2 (SOC2)

Qubole has completed an accredited attestation of the Trust Services Principles defined by the American Institute of Certified Public Accountants, pursuant to Attestation Statement 101 (AT 101). also commonly referred to as the SOC2. You can find our blog announcing this here. Qubole, like all other SOC2-compliant entities, requires a signed Non-Disclosure Agreement in order to share a copy of this report with a third party. Please contact your account executive if you are interested in obtaining a copy of the report.

Privacy Shield

Qubole has engaged with TRUSTe to complete and attest to compliance with the US Privacy Shield regulation around privacy and transfer of EU Personal Data to the United States.

ISO-27001

Qubole believes that compliance must be built throughout all components of our service. To attest to this commitment, our processes, procedures, controls, operations and activities align with the ISO-27001 standards, and are reflected in our policies and other attestation and certification work, including the Qubole SOC2 Type 2 attestation report. Qubole plans to certify ISO-27001 and ISO-27017 within the next twelve months.

Authentication / Authorization

Provisioning a Qubole account is easy: the Qubole Data Service provides several options for authenticating and provisioning users, including SAML, OAuth, ADFS and Whitelisting, as well as a number of granular access controls to ensure your users have the desired access privileges.

Creating an Account and Authorizing Users

To create a Qubole account, go to https://us.qubole.com and visit the Sign-Up link in the lower left hand corner. For security reasons, you will be sent an email to confirm and activate your account.

Single Sign On

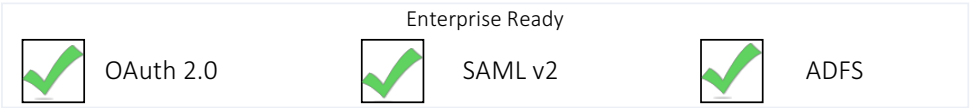

Unlike other vendors, Qubole doesn't charge you to set up the enterprise authentication and authorization scheme that best supports your business needs. Qubole supports all industry standard authentication and authorization protocols including OAuth 2.0, SAML 2.0 and ADFS. You may also freely use a number of user-management options to administer your Qubole users by navigating to Control Panel > Manage Users. To use any of these enterprise features please reach out to your account manager for assistance.

OAuth

Open Authentication is one of the authentication methods that can be configured for use with Qubole. OAuth essentially provides an application (a.k.a. *Consumer*) with a token validated against a trusted

Authentication provider, such as Google. Qubole supports Google OAuth 2.0 and Microsoft OAuth 2.0. Consult our reference documentation for more detailed setup instructions.

SAML

Security Assertion Markup Language is an enterprise authentication and authorization scheme 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 a number of third parties including Okta, OneLogin, Ping Identity and others who support SAML v2.0.

ADFS

Active Directory Federated Service allows an enterprise to connect to a corporate Active Directory to enable authentication that aligns with existing corporate requirements.

Configuring Roles

Qubole provides user-administration capabilities that allow up to 14 different resource types and functionality to be configured for individual users. For example, you can assign permissions to your data analysts and data scientists to create, run and modify queries and commands but not edit a cluster's configuration. Common concerns are addressed, such as preventing users from having rights to modify or affect the status of clusters, limit what commands they can execute, which data engines they can utilize and many more.

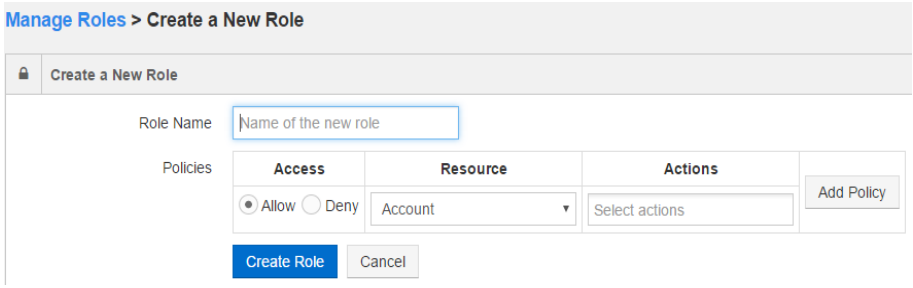

Per-User API Tokens

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

You can find your (user) authentication token by going to Control Panel > My Accounts > Show API Token

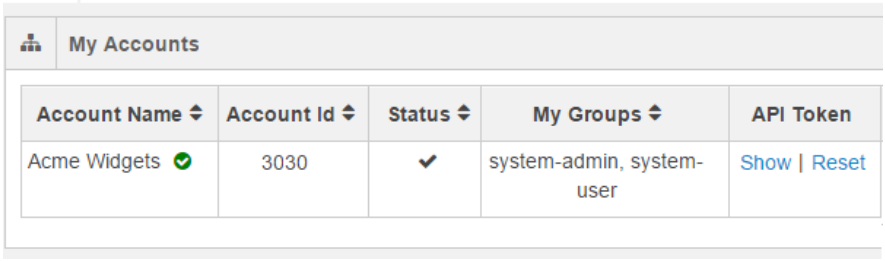

Accessing Data

Qubole takes advantages of AWS IAM roles to limit access to resources such as storage and compute. Using a refined set of permissions our customers to use Qubole on their behalf.

Using IAM Roles

Qubole uses the AWS Identity and Access Management (IAM) to limit the privileges available to each user to utilize the Qubole Data Service. Using IAM roles allows Qubole customers to limit the privileges granted to Qubole to just allowing the initiation and termination of instances.

Those instances are then formed into clusters that you define and configure. You can then grant a secondary role that allows Qubole to attach a data policy to those resources. This secondary or (dual) role only operates within the customer's account and is not accessible outside of that account by anyone including Qubole.

Qubole will initiate and terminate instances on your behalf based using the configuration which you create through the QDS control panel. For instance, you can set a minimum and maximum cluster size, define whether a cluster should utilize AWS Spot or On-demand instances, or a

combination of both, even defining a fallback strategy when SPOT instances are not available, and so on.

Data Encryption

Encryption Supported

Qubole supports end-to-end encryption throughout the Qubole Data Service (QDS) as defined below. Encryption protects the confidentiality of Customer Data ensuring that only certain users and entities can access specific data sets (please note that encryption may impact performance of QDS).

Data at Rest

Customers can configure compute nodes to use encryption when copying data to slave nodes, and reporting results to the master node through the Qubole UI. The latter can be found under the control panel inside the clusters configuration drop down.

Qubole also optionally supports AWS S3 Server Side Encryption, which ensures that when data is written to S3, it will immediately be encrypted on disk. This is configured within the Amazon AWS Console for the customer's account and is granted to Qubole via the IAM role specified above. Please refer to the AWS documentation found here and the Qubole documentation found here.

Using Customer Supplied Keys

Qubole also supports Amazon AWS KMS which stores keys in a shared Hardware Security Module dedicated for encryption key storage. For more information on how to configure and use this please refer to our documentation.

Data in Transit

Qubole supports Transport Layer Security (TLS) v1.2 (sometimes referred to as SSL or HTTPS) for communication between Customer browsers and clients, REST API endpoints and the Qubole Data Service. Qubole has enabled this by default.

Security of the Platform

Firewalls via Security Groups

Qubole requires the use of security in the creation of and maintenance of customer clusters. This is done at a minimum with secure permissions granting access from the Qubole Control Plane to the Data Plane which resides in the Customer's AWS account. All communication between the Control Plane and the Data Planes is encrypted. Qubole also uses these techniques to secure the various tiers of the QDS Control Plane.

Security Event & 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 certain variables exceed given thresholds, that event generated is forwarded to an on-call engineer for evaluation and escalation.

Intrusion Detection

Qubole deploys a log-based intrusion detection system that identifies anomalies on instances within the Qubole Control Plane. When an unusual event is detected in system or audit logs, this triggers alarms based on industry best practices which are forwarded to the SEIM (described above).

Qubole supports customer supplied and operated Intrusion Detection (IDS) or Intrusion Prevention (IPS) services within the customer account via bootstrapping scripts. Your Qubole Solutions Architect can explain how this works.

File Integrity

Qubole operates a file integrity (FIM) management service that detects changes in certain critical files on instances within the Qubole Control Plane such as changes to critical logs, Qubole application code, set UID binaries and additional configuration files and objects that may indicate a

host has been compromised. When an unusual file event is detected an alert is forwarded to the Qubole SEIM (described above).

Vulnerability Scanning & Detection

Qubole scans external systems to determine if there are any application or system specific risks, and works with the Qubole operations team to remediate these risks. Vulnerability detection looks for known software vulnerabilities to packages, applications and services used to provide the Qubole Data Service. These scans are configured to run weekly and are remediated based on severity.

Hardened Baselines & Configuration Standards

Qubole evaluates the baseline of systems deployed in the Qubole Control Plane on a weekly basis, using the Center for Internet Security (CIS) benchmarks. These benchmarks cover not only applications installed, configured and running on the systems, but also system and configuration files, including restrictions around permissions, logs, system binaries. This benchmark covers more than 150 discrete changes and is applied to all instances deployed to both the Qubole Control Plane as well as instances launched in Customer's accounts via the Qubole Data Plane Amazon Machine Instance (AMI).