At Databricks, we know that data is one of your most valuable assets and always has to be protected. That's why, in the light of the wave of recent cyber attacks involving unauthorized access, we wanted to take this opportunity to reinforce **the top 5 security best practices** that customers can leverage to protect their Databricks Data Intelligence Platform against this kind of attack:

1. **Use single sign-on (SSO) for *all* user access:** For Azure Databricks and Databricks on Google Cloud, the platform is natively integrated with Azure Entra ID and Google Cloud Identity, so you get this control for free. For AWS customers, we recommend that all customers configure single sign-on (SSO) between your identity provider (IdP) and your Databricks account via SAML 2.0 or OpenID Connect (OIDC). Although SSO can be configured at both the account and workspace level, Databricks recommends using a single account-level SSO configuration with unified login enabled for all workspaces.
2. **Leverage multi-factor authentication for *all* user access:** With Azure Databricks and Databricks on Google Cloud you can configure multi-factor authentication (MFA) via Microsoft Entra ID Conditional Access and Google Cloud Identity respectively. For AWS customers, we recommend that they configure MFA within their own identity providers. For all clouds Databricks advocates that MFA is enforced for all user access, including to the account console and even for emergency access scenarios. Where possible we also encourage the use of physical authentication tokens such as FIDO2 keys. These keys augment traditional MFA protections by requiring interaction with a physical token that cannot be compromised.
3. **Restrict access using IP access lists and/or PrivateLink:** IP access lists (AWS, Azure, GCP) restrict the IP addresses that can be used to access Databricks by checking if the user or client is coming from a trusted IP address range such as a VPN or office network. Established user sessions do not work if the user moves to a bad IP address, such as when disconnecting from the VPN. Databricks recommends that customers configure IP access lists for their Databricks account (AWS, Azure, GCP), workspaces (AWS, Azure, GCP). and Delta Sharing recipients (AWS, Azure, GCP). Customers that require private networking can leverage Private Link on all three clouds.
4. **Protect tokens like you would protect credentials:** Where possible customers should use OAuth or Microsoft Entra ID tokens for authentication. OAuth reduces risk by using short-lived (one hour) access tokens which are automatically generated and refreshed without needing to be embedded in code or configuration. If Personal Access Tokens are not needed, Databricks recommends that customers disable them. Where Personal Access Tokens are required, Databricks recommends that customers control who can create & use them, set a short maximum lifetime and monitor and revoke them when necessary. Please see the documentation for AWS, Azure and GCP for more information.
5. **Configure audit logs and monitor for potential Indicators of Compromise (IoCs):** Audit logs serve as your system of record for all of the material events happening on your Databricks Data Intelligence Platform. Databricks provides several options for customers to leverage their audit logs, from built-in querying via System Tables to delivery to the customer's cloud account. Databricks recommends that customers configure and monitor their audit log for potential Indicators of

Compromise. Please refer to this blog and companion GitHub repo for a set of out-of-the-box examples to help get you started.

The best practices above are the **5 most important protections** that customers can leverage to protect their Databricks Data Intelligence Platform against unauthorized access. But for our most security-conscious customers there are plenty more to choose from! You can find out detailed security best practice guides for your chosen cloud(s) on our Security & Trust Center. If reading isn't your thing, you can also leverage our Security Analysis Tool to perform an automated security health check and simply review the results.