Security at Colabra

Colabra

A modern cloud based security approach

Introduction

As cloud services have evolved, new approaches to software architectures in general and security architectures in particular have been available to software vendor organizations of all sizes. IaaS, PaaS, and more fine grained services provide options that were unavailable to all but the largest organizations up until the recent past. This commoditization of capabilities and expertise, specifically security capabilities and expertise, frees up vendors of all sizes to focus on their core competencies while maintaining the highest levels of security and compliance. This is possible by allowing organizations to effect the following:

- Transfer of attack surface risk
- Strategic deployment of security resources
- Leveraging third party security expertise
- Compliance and best practices based security implementation
- Transparent response and remediation

Transfer of attack surface risk

One of the overarching goals of security remains to reduce the size of the overall attack surface. However, the cloud services options mentioned above allow the attack surface of the commodity aspects of the platforms/services to have their risk transferred.

Risk transfer is a longstanding and accepted practice for addressing security risks. Vendors choose to transfer risk when that risk can be best addressed by third parties, if the internal addressing of that risk would be cost prohibitive, or if there are contractual or regulatory requirements that can best be met by the use of a third party.

Risk transfer is feasible through legal mechanisms including, but not limited to, third party certifications, third parties answering security questionnaires, data processing agreements, contractual obligations, etc. These mechanisms may only be limited to the parties involved through contracts, be established in accordance with applicable regulations, be part of accepted industry compliance practices etc. Organizations choose which method or methods best meet their requirements and are in accordance with their business and security goals.

Leveraging third party security expertise

Due to the scope of their operations and their compliance requirements, cloud service providers typically maintain specialized security teams. These teams are focused on the unique security requirements of the platform/services their organization provides. As such, they have the expertise and the resources to best prevent or detect and remediate security incidents involving the provided platform/services. As such, cloud service providers have a mandate to maintain the Confidentiality, Integrity, and Availability of their provided platform/services. This allows organizations to leverage all of the above to the advantage of their unique security posture.

Strategic deployment of security resources

All of the above allows organizations to focus their security resources on their own unique codebase. No matter the size of an organization, security budgets are always subject to their business goals and acceptable levels of risk. This is especially the case with startups and smaller organizations. Given the increasingly complex regulatory environment, growing customer awareness of security incidents, and burgeoning penalties for security breaches, the security obligations of organizations are becoming more and more pronounced. As such, organizations need to maximize the value of their, often limited, security resources. The best way to do so is by focusing entirely on hardening their own codebase.

Compliance based security implementation

The described approach to security requires an implementation that is compliant with applicable regulations and industry standards as required by the organization's business goals.

While it may be tempting to take a narrow interpretation of "applicable regulations", the fact of the matter is that modern security regulations are often drafted to be as widely applicable as possible, including applicability beyond their national or regional boundaries. This is especially complex due to the connected nature of individuals, organizations, and nations in the modern digital world.

Examples of such regulations typically include the GDPR, the CCPA, and the CLOUD Act. In addition to regulations, compliance with industry standards such as PCI-DSS and SOC 2 may be required. Standards such as PCI-DSS are prescriptive, i.e., they mandate certain aspects of implementation and standards such as SOC 2 are descriptive, i.e., verifying that an organization's chosen implementation is as described.

In the context of this approach to this Modern Cloud Based Security Approach, a minimum set of controls will be delineated. These controls will be implemented with respect to an organization's unique code base; i.e., not including those required to address the security risks transferred to third parties.

Controls

Two-person integrity

A control that requires at least two people to take actions that affect security. Such actions may include:

- Merging commits to the production branch of source control
- Deployment to production infrastructure of new product/service versions
- Granting and revoking access to sensitive resources

Principle of least privilege

A control that ensures no person has more privileges than required to do the tasks in their job description. A simple example would be that a non-administrative developer shouldn't have the privileges required to onboard and offboard other development staff.

Access control lists

A control that ensures that a canonical list is maintained that shows what privileges are required to gain access to secured resources and who has those privileges. A comprehensive description of access control techniques is beyond the scope of this document, there are many different approaches to this topic. N.B.: these different approaches are sometimes described in different terms than "Access Control Lists", but that term is used here for simplicity.

Audit trails

A control that records access to secured resources. These may take the form of normal system event logs, custom logging tools, etc. It's often necessary to retain these logs for a required and/or specified time period.

Internal audits

A control that utilizes automated or manual reviews of the status of other controls, audit trails, seeks to identify potential security incidents, etc. These may take place on an scheduled basis, an ad hoc basis, or both. When they take place may be required by regulations, industry standards, and/or contractual obligations.

Automated source code analysis

A control that identifies stylistic issues that may degrade the development process, common coding mistakes that otherwise pass the normal syntactic and semantic analysis stages, security vulnerabilities, etc. While this control may be run manually as mandated by policy, it is better utilized as part of the process of committing or merging changes in version control, or as part of CI/CD regime.

Automated security analysis

A control that seeks to find any security issues not found by the other controls. Controls if this class may take many forms involving many different automated tools. This analysis may be run on a scheduled basis, an ad hoc basis, or both. When it takes place may be required by regulations, industry standards, and/or contractual obligations.

Addendum: Transparent response and remediation

Any description of an approach to security would be incomplete without a section describing what should take place in the event of a security incident. The below describes the minimum set of actions that should take place and what requirements they should be subject to. However, as in all such cases, the below should be evaluated in the context of regulatory requirements, industry standard requirements, and contractual obligations.

Notification

As soon as the incident response process has been completed to the point where it can be determined that a breach has taken place and who has been affected, the affected parties should be notified as soon as possible. However, the time frame of the notification, what information is disclosed, and if third parties must also be notified may be governed by applicable regulations, internal policies, or contractual obligations. Notification channels should be pre-established, have backups, and be able to be executed immediately.

Investigation

Whether part of or independent from the incident response process, an investigation needs to be initiated to determine the extent of the damage to clients. This is separate from any other damage that may have been done to the organization or its data. This is purely for the purposes of assessing client impact and to assist in any future remediation efforts whether they be initiated by the organization or the client.

The investigation should also determine if the incident occured for the express purpose of targeting the client. If that is the determination, the investigation should be expanded, within the legal capabilities of the organization, to the identification of bad actors, including the possibility of rogue employees.

Cooperation

In the event that a security incident leads to a prosecution or to civil litigation of a third party on behalf of a client, the organization should offer assistance to any law enforcement bodies or the affected parties' legal teams. Unless there is a compelling reason or obligation not to do so, the organization should act proactively to provide any statements, testimony, data, etc that would provide assistance to the prosecution or civil litigation.

MORE INFORMATION

We are confident that Colabra can meet your lab's workflow and security needs.
Please contact us at trust@colabra.app with your questions about data privacy and security.