TECHNICAL BRIEF | December 2015

CA Risk Authentication delivers Intelligent Risk-Based Authentication for the Enterprise

Robert Marti
Product Marketing, CA Security CSU

ca technologies

Table of Contents

Executive Summary

Challenge

Passwords are often a critical weak link in a security system, and they fail to satisfy many industry best practices and regulatory guidelines for protecting identities and data. In addition, many other types of login credentials have also been stolen and used to breach the enterprise. Today's approach of relying solely on a credential to validate user identity has shown to be vulnerable to attacks. Organizations need a more secure way to protect access to corporate applications and data and protect their users from account takeover. Combining intelligent, contextual-based risk analytics with step-up authentication has proven an effective approach to combating online identity fraud.

Opportunity

User behavior is a strong indicator of an individual's identity. CA Risk Authentication provides a secure, user-convenient and cost-effective way to protect sensitive applications and data. CA Risk Authentication can detect and protect against high-risk login attempts or other sensitive access by analyzing a wide set of factors, including user behavior, device characteristics, geolocation and velocity data, without requiring any direct input from the end user. And when deemed too risky, the user can automatically be prompted to submit additional credentials or information to further prove their identity.

Benefits

CA Risk Authentication provides additional security for any application or data that can be accessed through a browser or mobile device and thus helps reduce the risk of online identity fraud and other data breaches. Simply adding risk analysis provides greater assurance that the user is who they claim to be. Providing this protection in a transparent manner allows organizations to improve security and reduce fraudulent access without impeding the productivity of employees or partners or impacting customers. In fact, CA Risk Authentication provides an additional opportunity to improve the user experience by limiting authentication only to the times judged risky, which matches user intuition.

Section 1: Risk Analysis Fundamentals

Authentication That Is Convenient, Secure and Cost-Effective

In this mobile world, information, applications and employees no longer stay inside the enterprise network but work from wherever the business needs them. Many companies pursue the mantra "Access from anywhere, anytime, any device." Security must exist in this world of disappearing perimeters. Organizations must both embrace and secure the open enterprise.

In addition, your customers are far more likely to experience your brand and interact with your enterprise through software than a live person. To thrive in this new reality, every enterprise needs to have a mobile strategy and online presence. However, these applications are also increasingly exposing financial, healthcare, and personally identifiable data, and new breaches are being reported daily in the news.

Stronger authentication mechanisms are rapidly becoming a necessity for all Internet-facing and mobile-based applications. Contextual, risk-based authentication provides a smart, zero friction approach that addresses security concerns without inconveniencing users. It learns individual users' behaviors and detects when they are deviating from their normal patterns, which makes it more personal and more accurate than a static credential.

CA Risk Authentication

CA Risk Authentication provides a robust, multi-channel, configurable risk assessment engine, which can assess the risk for a specific transaction based on four categories: Device Identification, Location, User Behavior and Velocity. The risk evaluation produces a risk score that is combined with business policies to determine if any action is required. This allows legitimate users, most access attempts, to continue uninterrupted because their risk score is low. In the case of an unacceptable risk score, where the user's behavior is outside the norm, the user can be required to do a step-up authentication to further prove their identity. For example, Mark works from the office in New York and rarely travels for business, but suddenly he is asked to present at a conference in London. When he logs in from his London hotel, the system recognizes that this is unusual behavior for Mark and prompts him for additional login information.

Sophistication of fraud continues to increase making it more difficult to distinguish between fraudulent and genuine users. As enterprises expand their applications beyond the perimeter they need more than "black box" risk assessments. The ability to define risk rules for any transaction, control risk parameters on-the-fly, and understand fraudulent and legitimate behavior is essential in any authentication deployment. CA Risk Authentication provides:

Better User Experience. Can authenticate an employee, partner or customer without complicated user credentials. Enterprises can use risk-based authentication with simple passwords and only require step up authentication if the login seems risky. This provides a frictionless login experience for most users improving satisfaction and/or productivity while reducing calls to the customer service representatives or help desk.

Greater Accuracy. Using an enterprise-specific model that understands legitimate and fraudulent behavior we can determine the validity of a user in context of what is normal for that individual. In real-time during authentication, CA Risk Authentication takes a multi-dimensional view of the login by using elements such as the device characteristics, geo-location, login velocity and historical user behavior.

Faster Speed of Change. CA Risk Authentication allows enterprises to make changes to rules and business policies on the fly. With access to case information for all logins, an administrator can update authentication policies based on an action that triggered step-up authentication or denied login. Policies can be updated or added dynamically as required by the business needs. There is no dependency on the vendor to make changes. In addition, the user behavior model learns over time. As user behavior patterns change, the model automatically updates itself.

Elements of Risk Analysis

The elements of an intelligent contextual authentication solution includes behaviors, device and location. Each is described in greater detail in the following sections.

Figure 1.

Contextual risk-based authentication is being added to traditional authentication credential types in order to create a layered security approach to prevent unauthorized access and data breaches.

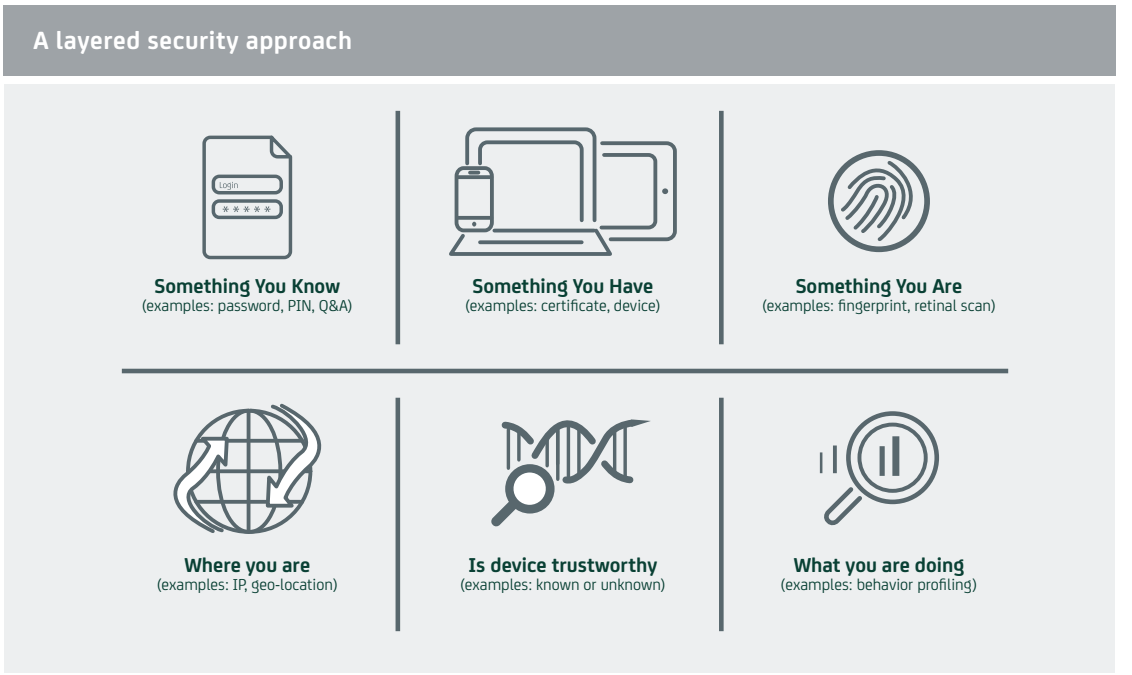

Device Identification

CA Risk Authentication provides a patented approach to device identification, called DeviceDNA™, which consists of two separate components:

- **Device ID.** This is a unique identifier that we assign to the device. The Device ID is can be stored on the device using commonly accepted browser technologies or a reverse lookup can be done with matching information stored on the server.
- **Machine Fingerprint**. This is a set of characteristics we observe from the device that allow us to verify the authenticity of its Device ID or find it using a reverse lookup process. It includes browser, machine and software data, and a number of parameters about the device type itself.

The Device ID is our method to "name" the device and the fingerprint is our method to validate the device. Using the Device ID, we create associations between the devices and the users in a many-to-many mapping. We also store a history of devices used in our database along with the prior fingerprints of each device; these are stored such that personal user information cannot be identified for privacy concerns. This network of associations allows us to make a number of judgments about both devices and users. For example, if the user has established a history of using this device, it can be considered a second factor of authentication (something you have), which increases our confidence in the user's identity.

Current technologies used for gathering fingerprint data are the HTTP header protocol and JavaScript collectors for browser-based transactions. Customers actively use this capability to collect DeviceDNA from various devices including tablets, smartphones, televisions and other non-pc devices. The solution also provides Mobile DeviceDNA libraries, which can be embedded into first-party mobile apps in order to gather data for a risk assessment from the mobile device.

Location

CA Risk Authentication allows administrators to define geo-location risk rules that can evaluate the data collected from the end user's browser or mobile device. The location information is derived from the end user's IP address by leveraging third-party Neustar IP Intelligence data, which supplies several kinds of data, including:

- **Geo-location data.** This classifies each IP by Latitude, Longitude, Continent, Country and City. The data is used in the mobility calculations and negative country checks. The data can also be used in add-on rules for a number of other checks. The Continent field (Africa, Antarctica, Asia, Australia, Europe, North America, Oceania and South America) and the City field (approximately 95,000 cities) are both useful for courser or finer granularity checks on the point of access.

- **Connection information.** Each IP is classified by routing type, connection type and line speed. This information, especially Routing Type, is useful in assessing the validity of the geo-location information. For instance, if the connection type is Satellite then the user's location is not reliable. However, "Fixed" connection types (e.g., Cable or DSL) are less likely to be origins of fraud since their locations are more easily back-traced to internet accounts and that therefore this data can be valuable in fighting fraud.
- **Anonymizer data.** A third category called anonymizer data is also supplied through Neustar, which performs rigorous testing of IPs to determine if their location information is reliable. As part of this they identify IPs as "Anonymizers." IPs with this status have tested positive as anonymous proxies that are used to hide the true location of the end user. While this does not necessarily indicate that the intent is fraudulent, it does clearly indicate that the user is hiding their location and thus represents a higher access risk potential.

The Neustar IP GeoLocation database updates are included as part of the CA Risk Authentication software license and annual maintenance.

User Behavior

CA Risk Authentication includes an enterprise modeling technique called User Behavior Profiling, which learns individual user behavior and allows customers to do step-up authentication when their end user's behavioral patterns deviates from their norm. This feature adds another dimension to the risk authentication process by performing a statistical assessment of the user's current observed behavior against their established personal history.

Observed behavior will form a consistent pattern over time that is user-specific. For example, some users will exclusively login from one location; others will be observed to commute or travel. Conventional systems will tag certain locations with scores that rate the risk of access from each city based on expert opinion or training data. This approach can be effective when such data or opinion exists, but it creates a dependency on that data. On the other hand, our approach to does not have this dependence. It centers on the user, not the data. In this case, our system doesn't judge the goodness or badness of a location; it simply scores how normal or abnormal that location is for each particular user. Thus, the user's pattern of access is important, not the location(s) themselves.

With CA Risk Authentication User Behavior Profiling, assessments are performed across several dimensions and combined into one model score that measures how consistent the current observation is to that user's history. Specifically, the solution evaluates the following types of data:

Table 1.

User behavior profiling represents a new generation of intelligent risk analytic tools.

User behavior profiling data elements

Observed Data	What Does It Tell Us?
Action choices	What does the user intend to accomplish? Is this action consistent with previous action history?
Browser choices	Does the user demonstrate a preference for a particular browser type or browser language, or does he regularly utilize several types or multi-lingual?
Device choices	Does user own certain devices and change only when they purchase a new phone, tablet or computer? Does user tend to use one device or many?
IP choices	Does user tend to stick with the same small set of IP providers and connection types, or does user tend to use more new IP providers than others?
Location choices	Does the user travel frequently or mostly stay at home? Some users are extremely predictable in their location of choice while others travel extensively.
Time choices	Does the user exhibit a specific preference for time of day or day of week pattern or not? If so, this can be a good predictor of future behavior.

User Behavior Profiling is an effective detector of many threats and fraud attacks. Not only must the attacker understand and mimic the true user's behavior, they also must constrain themselves to performing work that the user normally performs. In addition, it also provides the following benefits:

- **No training required.** Can be deployed in production without collecting prior history. Some systems require training on a sample data set that is generally difficult to obtain and often contains PII data.
- **Quick time to value.** Starts yielding results on the second observation. Since User Behavior Profiling calculates normalcy, the assessment of even the second transaction is relevant.
- **Automatically adjusts over time.** There is no ongoing customer analysis of good/bad transactions. As the user's behavior changes, the model automatically learns the new patterns.
- **Intuitive user experience.** Users are not surprised if they are required to (re)authenticate when they do something that may be deemed unusual and see this as good security practice.
- **Difficult to reverse engineer.** The combination of multiple assessment dimensions makes it extremely difficult to reverse engineer because behavior is complex and accumulated over time.

Risk Analysis Process

CA Risk Authentication can "score" any transaction at any point during the application session (e.g., before, during or after authentication) where a transaction can be anything defined by the customer (e.g., registration, login, access request, funds transfer, purchase, download, etc.). The risk evaluation process begins when a user attempts to perform one of these defined transactions. In addition, the

Figure 2.

Risk analysis can detect and stop inappropriate access on its own to catch fraudulent activity even if credentials have been compromised.

Risk analysis process flow

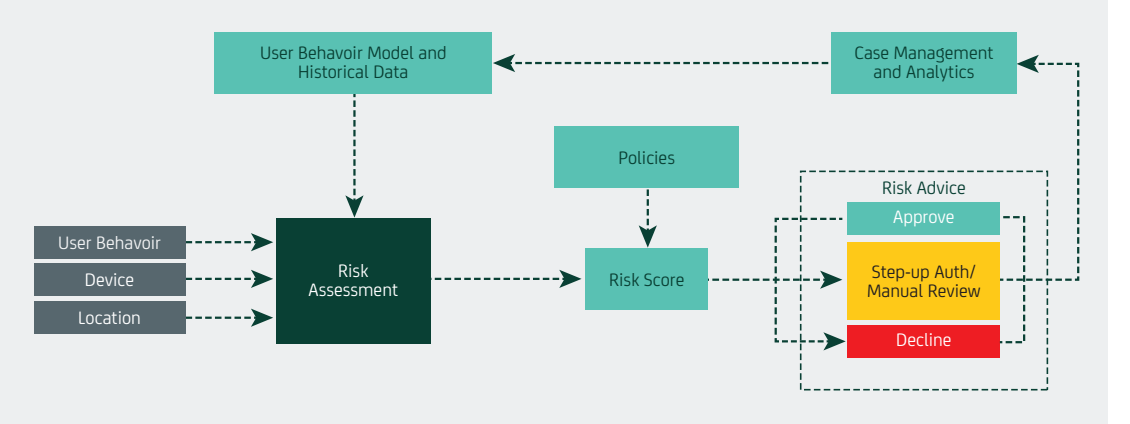

solution allows organizations to define a different set of risk rules and policies for different transactions, different devices and/or access channels, and different user communities. Data is collected and forwarded to the risk engine, which assesses risk using the following process.

As shown, a risk score is derived from the risk assessment, which consists of two parts. First the enterprise model will analyze the data to determine if the user's current behavior is deviating from their normal usage patterns. This result is then forwarded to the rules engine, which analyzes this data along with the results from the rest of the risk rules that have been defined for this transaction. The outcome of this analysis results in a risk score, and based on predefined policies, the risk advice will result in one of the following actions:

- **Approve the Transaction.** In this case, there is a high degree of confidence in the user's identity and other parameters of the requested action look to be of acceptable risk, so no further action is required. An "approved" message is returned to the application, which can then process the transaction normally.
- **Alert.** In this case, while there is confidence in the user's identity, this request merits follow-up examination by the help desk or security team. An "approved" message is returned to the application; however, an email is sent to an administrator, so that this transaction can be investigated at a later time to determine if it was fraudulent or not. Additionally, an alert can be a trigger to the enterprise security control center.
- **Initiate Action.** In this case, the risk score was above the policy threshold, but can be remediated, so the risk engine initiates one of the following actions:
  - **Step-Up Authentication.** The solution prompts the user to submit additional proof of identity before the transaction will be processed. If the user successfully completes the step-up authentication, the solution will return an "approved" message to the calling application, which could then process the transaction normally.
  - **Manual Review.** The transaction can be forwarded to the case management console for review and disposition by a Fraud Analyst or Customer Service Representative. This representative may contact the user to determine if the transaction is legitimate or fraudulent.

- **Reject the Transaction.** In this case, the risk score is above a policy threshold, and this is deemed a high-risk transaction. The risk engine would return a "declined" message to the application, which would then alert the user that the transaction cannot be processed and that they should contact customer support.

This array of mitigation choices allows you much more flexibility in how you manage your risk–creating a balance between user convenience and security that is appropriate for your organization.

Preventing Data Breaches with Risk-Aware Session Management

When users engage with your organization, authentication is just the first step. How do you continue to protect the user throughout their ongoing activities within your applications? Session hijacking has been identified as one of the top security risks in the recent OWASP survey that was published in 2013, and for good reason—what happens during a session matters. In many cases, it is sufficient to simply authenticate a user, and with that process alone you have validated that the user has a legitimate reason to access the site. However, when the user begins to actively engage with the site, especially when accessing sensitive personal or financial information, stronger security may be required.

Figure 3.

Risk analysis can help establish a virtually impenetrable shield between your users and the bad guys looking to hijack their sessions.

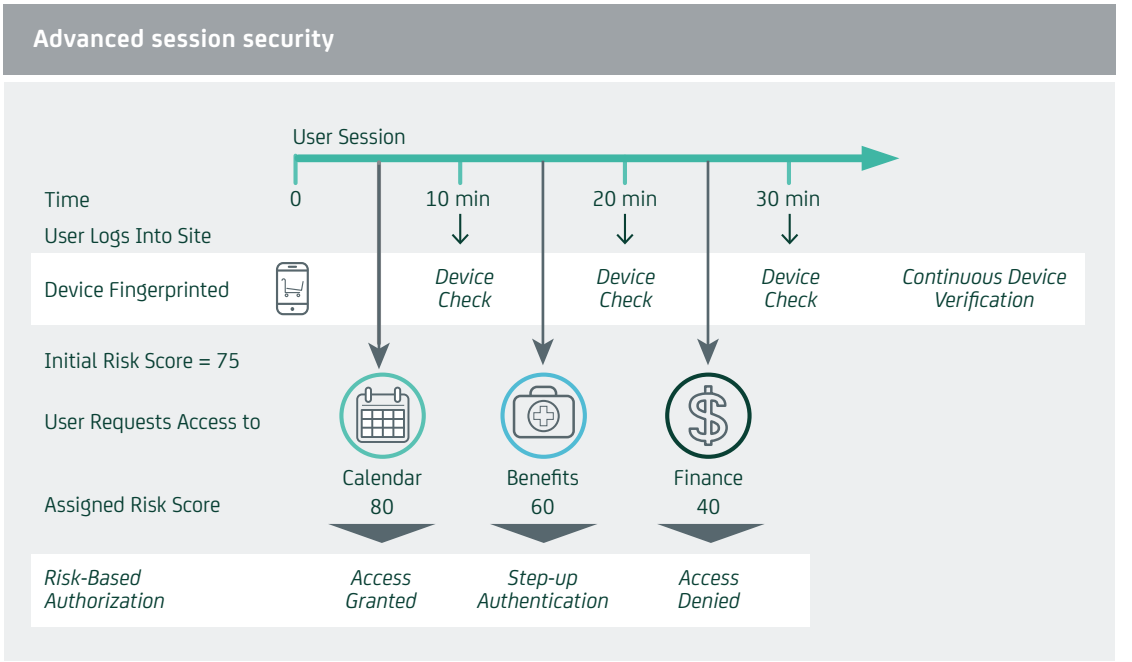

CA Risk Authentication has been integrated with CA Single Sign-On to combat session hijacking via two approaches:

- **Continuous device verification.** Using this method, the solution reconfirms that the user who initiated the session is still in control. This is accomplished by fingerprinting the device at login, and then pinging the device at predefined intervals, so long as the session remains open, to verify whether the same device is connected or not. And if the device connected to the session fails to verify with the proper identifier, the session is terminated and the connection severed.
- **Risk-based authorization.** Using risk-based authorization, each transaction is assigned a specific risk score threshold. When the user logs in, a risk score is generated. This score may not prevent the user from accessing the site. However, as the user attempts to access more sensitive areas of the application/site, the risk score can be evaluated to determine if it is above or below the threshold. Whenever the user's risk score fails to qualify for the location they are trying to access, more rigorous step-up authentication can be required or the user can be denied access.

Each of these can be extremely helpful in securing user sessions on their own, but when used together, they can help establish a layered security approach to provide advanced session security.

Configurable Rules Engine

CA Risk Authentication was designed using a "White Box" philosophy by providing a configurable rules engine, which allows administrators to create and/or modify risk rules on demand to balance user convenience, threat mitigation and back-office costs. In addition, to support initial deployments, the solution provides a standard set of default rule types, which can be used to evaluate a wide range of transaction and session data.

The following table briefly describes the default rules and recommended starting points for risk advice. These can be changed or modified by the enterprise based on their specific needs and requirements.

Table 2.

Best practice rule sets are provided that cover typical fraud patterns.

Default set of risk rules

Rule Name	Recommended Risk Advice	Usage
Device-Based Rules		
Browser Type	Fixed Input	This rule tests for the type of browser being used for access. This data is generally used for device fingerprinting; however, you can also configure specific rules that can create an alert, require step-up authentication or deny access when the user is coming in from a specific browser type.
DeviceID Check	INCRAUTH	This rule will check to see if the device has been previously tagged with a Device ID, and is a known device. When a user is coming in from a new device, this rule initiates a step-up authentication.
Device Type	Fixed Input	This rule tests for the type of device being used for access. This data is generally used for device fingerprinting; however, you can also configure specific rules that can create an alert, require step-up authentication, or deny access when the user is coming in from a specific browser type.

Table 2.

Best practice rule sets are provided that cover typical fraud patterns.

Default set of risk rules

Rule Name	Recommended Risk Advice	Usage
Device-Based Rules		
Device Velocity	ALERT	This rule tests for the number of successful or attempted uses of the device, and creates an alert for any device attempting access more frequently than the threshold you define.
Device Velocity with Multiple Accounts	ALERT	This rule tests for the number of different accounts that have been accessed or attempted to be accessed within a specified time, and creates an alert for any device attempting access more accounts than the threshold you define.
Machine Fingerprint (MFP) Match	INCRAUTH or DENY	This rule tests to determine if the current device fingerprint is a sufficient match to its previous fingerprint, and can require step-up authentication or deny access if the match percentage is below a defined threshold.
Negative Device Check	DENY	This rule tests to determine if the device has been placed on a black or watch list, and can deny access to any devices found on these lists.
OS Type	Fixed Input	This rule tests for the type of device OS being used for access. This data is generally used for device fingerprinting; however, you can also configure specific rules that can create an alert, require step-up authentication, or deny access when the user is coming in from a specific OS type.
User Associated with Device and MFP Matched	ALLOW	This rule tests if the user has an established a successful history (in time or number of transactions) with a device, and will allow access when the user is using a known and established device.
User Not Associated with Device and MFP Match	INCRAUTH	This rule requires additional authentication when the user comes in from a known device, but one that has not been previously associated to the user.
Location-Based Rules		
Negative Country Rule	DENY	This rule tests for the geographic region from which the access request is originating, and can deny access to users attempting to access from any countries that have been blacklisted. This rule tests for the geographic region from which the access request is originating, and can deny access to users attempting to access from any countries that have been blacklisted.
Negative IP Check	DENY	The Negative IP Check Rule performs two functions within a single rule. First, it checks the end user's IP against the list of known anonymizer proxies. IP addresses are classed with an anonymizer status. The definitions of this status is: Active; Suspect; Private; Inactive; and Unknown. You can control which kinds of anonymizer IPs you include in this rule. Secondly, the rule consults the Negative IP Address List that you define to verify whether the IP is in one of the ranges defined in your negative IP black list. The solution provides a wizard that allows you to add single IP addresses or an entire range to this list.

Table 2.

Best practice rule sets are provided that cover typical fraud patterns.

Default set of risk rules

Rule Name	Recommended Risk Advice	Usage
Location-Based Rules		
Trusted IP Check	ALLOW	Similar to above, the solution allows you to "white list" single IP addresses or an entire range, including aggregators, via a wizard. This rule allows you to approve/allow access from the IP addresses on this list.
Zone Hopping	INCRAUTH	The location latitude and longitude are used in the Zone Hoping rule. This rule checks the required speed of physical travel required for the user to make successive transactions from the IPs they used for access. If they are "traveling" too fast then we must conclude that either two people were accessing the account or the user did something, either intentionally or inadvertently, to mask their true location.
User-Based Rules		
Behavior Model Score	INCRAUTH or DENY	This rule tests the resultant user behavior model score to determine if it exceeds the defined threshold, and can require step-up authentication or deny access based on your preference.
Exception User Check	ALLOW	Regardless of how good your rules are constructed you may occasionally block a good user from the system. For this reason you should always include an "override" that allows these users access, even though they are working under conditions that fall within the purview of some of your rules. This rule returns an "allow" risk advice for a user on the exception list regardless of the results from the other risk rules.
Unknown User Check	INCRAUTH	This rule require that all first time users perform an increased authentication.
User Velocity	ALERT	This rule creates an alert for any user attempting access more than frequently that a threshold you define. More frequent access may indicate a training issue or a penetration attempt.

The solution also provides a wizard that allows customers to modify the existing risk rules; build new rules; and define what actions to take when a risk score reaches a specific threshold. Evaluations rules can also be combined into different rule sets for different transaction types, devices and user groups. Rules can also be combined such that Rules A, B and C must all be true in order to trigger a "true" event.

In order to provide you with full data and information about your user's actions, we evaluate and report the results of all rules. These data are available in the reports and in the case management system. This information can be used to adjust your rules and policies to adjust for emerging threats or balance risk against user convenience. Finally, it should be noted that risk analysis can be run in "silent mode" to allow rule effectiveness to be evaluated without impacting users.

External Fraud Systems, Scoring Engines and Lists

CA Risk Authentication can also be integrated with external fraud and/or risk scoring systems to validate or augment its own risk assessment through two callouts:

- **Evaluation Callout.** This is executed as part of risk evaluation and is used to gather externally-based data to be used within the risk rules.
- **Scoring Callout.** This is executed after the standard scoring logic has executed and is used to pass risk scoring results to an external scoring engine.

Risk scores from multiple systems can also be aggregated to generate one combined score. For example, a user normally resides in Los Angeles may be logging in from New York—a suspect transaction. But the custom callout to a credit card authorization system may show "card present" transactions in New York that will confirm that the user is in New York and therefore reduce the risk of the online access.

Finally, CA Risk Authentication allows administrators to define risk rules that can check any input parameter against a user-defined list and then use that comparison outcome to assess risk, grant exception or deny access. These lists can be static or dynamic, and because you control the semantic meaning of the list, it can be a black list, grey list or white list. For convenience, the solution provides two common list-based rules for negative country and negative IP checking. However, most customers quickly build upon these to create more complex expressions.

Step-Up Authentication

CA Risk Authentication provides contextual adaptive authentication and can challenge a user for additional authentication credentials when their risk score exceeds a defined threshold. The following challenge methods are supported out-of-the-box:

- **Knowledge-Based Authentication.** The solution can support [X] number of challenge questions from a configurable list of known question/responses. This could be existing knowledge-based questions and responses or those captured by the solution.
- **Out-of-Band Authentication.** The solution can generate a random one-time password (OTP), which can be delivered to the user over email, text or voice channels.
- **Client-Side OTP.** Users can also generate an OTP on their mobile devices, using the CA Mobile OTP app, if the CA Strong Authentication product has been fully licensed and deployed in the environment.
- **Customer-Defined Authentication.** The solution can be customized to support any third-party credential for step-up authentication, such as fingerprint, wearable, out of wallet question, etc.

Section 2: Enterprise-Scale Capabilities

CA Risk Authentication includes several enterprise-scale features that ease deployment and ongoing administration. These capabilities are described in the following sections.

Auditing and Reporting

CA Risk Authentication stores all risk evaluation requests and their outcomes, as well as all administrative activity, including any changes to configuration, risk rule or user data. Audit data is written to an external relational database. In addition, the solution also exposes a rich set of information about each transaction within the case management interface as shown below.

Table 3.

Risk analysis reports enable organizations to understand why users are being challenged and adjust their rules as needed.

Sample risk analysis report data

Location Details			
IP Address	66.161.103.8	City	Washington
State	District of Columbia	Country	United States
Connection Type	Unknown	Line Speed	Unknown
IP Routing Type	None	Anonymizer Type	Unknown
Risk Assessment Details			
MFP Match %	91.66687	User Known	YES
Exception User Check	NO	Device MFP MAtch	YES
Trusted IP/ Aggregator Check	NO	Untrusted IP Check	NO
User Velocity Check	YES	Device ID KNn	YES
Device Velocity Check	NO	User Associated with Device ID	YES
Device Details			
Device Type	iPad	OS	IOS
Browser	Mobile Safari	Device ID Status	READ

CA Risk Authentication also provides standard risk analysis reports that include all risk analysis requests and recommended actions. These reports also include statistical summaries and detailed case analyses.

The solution also provides a series of reports that provide statistics and metrics on the system performance, including the following:

- **False Positives Report.** This report shows the transaction activity by rule annotated with fraud, not fraud and to be determined.
- **Fraud Statistics Report**. This report displays the overall statistical data for risk analysis results that the system generates in the specified time period.
- **Rule Effectiveness Report**. This report displays historical trends for rule activity.

These reports can be viewed on-screen through the Administrative Console or exported for further analysis.

Case Management

CA Risk Authentication provides case management capabilities. High risk transactions (cases) can be forwarded to a queue, where they can be reviewed and remediated via a web-based interface. The case management component provides a single unified view of the data associated with these cases. This enables administrators, fraud analysts, or customer service representatives to review, analyze, and remediate these high-risk transactions more efficiently. In addition, this interface can also be used to evaluate any transactions that caused an alert. In this case, the transaction was allowed to continue; however, because of the alert, the organization can now investigate these cases to determine if they were legitimate or not. The solution tracks the status and progress of all cases and maintains complete case histories with instant access to all related information. As cases are reviewed and determined to be True, False Positive or False Negative; this data can be used to modify the risk rules and policies.

High Availability

CA Risk Authentication solution is designed for high availability, including the following features:

- Stateless servers for deployments behind load balancers
- Distributed architecture for deployment across multiple geographic locations
- Connection pooling technology
- Transparent and automatic internal failover and failback capabilities
- Open standards and protocols
- Optimized database interface/schema
- Local memory caching
- Multi-CPU aware memory managers and false contention avoidance

Integration

CA Risk Authentication can easily be integrated with external applications and security systems via the following approaches:

- **JavaScript Library.** The solution provides a JavaScript DeviceDNA collection library for inclusion in your web pages.
- **Mobile Libraries.** The solution provides mobile libraries, which can be embedded into first party mobile apps in order to gather data for a risk assessment from the mobile device.
- **Web Services.** The solution provides RESTful APIs that can be used by external applications to submit risk data and request risk assessment.

CA Risk Authentication also provides an out-of-the-box adapter that allows the solution to be easily and quickly integrated with CA Single Sign-On. Similar adapters can also be built for other single sign-on and Web access management solutions.

Multi-tenancy

CA Risk Authentication is also designed to support multi-tenancy. The architecture allows for separate "organizations" for different user communities, applications, business units, etc., which might have different rules, risk profiles and usage characteristics. In addition, the administrative console is also constructed in a hierarchical style, which can delegate privilege and control to specific organizations to the appropriate level of administrator. Therefore, an administrator from one "organization" cannot view the users, risk rules (or results) or configuration data from another "organization."

Performance and Scalability

CA Risk Authentication was designed to provide the best performance possible, in order to meet the rigorous demands of our financial services customers, who are supporting millions of users. The three most important factors that impact performance are design, scalability and latency. The solution response times are a function of adequate hardware sizing to support peak loading (i.e., risk analysis requests per minute). CA provides hardware sizing guidelines within the Implementation Manual, and CA Services can also provide assistance during the architecture design phase of the deployment.

The second factor, scalability, addresses how well the solution provides services to thousands to millions of users while performing hundreds to thousands of authentications per minute. The solution is architected to scale horizontally and vertically to match your current architecture design with regard to load balancing, high availability, capacity, disaster recovery, etc. Vertical scaling is achieved through increasing memory, disk and processors. Horizontal scaling is achieved through additional local or remote server via load balancers, and provides performance gains, as well as high availability for critical deployments.

Finally, latency describes the delays introduced in the overall transaction time by the different components that make authentication service. The solution has been designed to minimize latency within its components; however, we have found that the largest contributor to total overall response time is network latency.

Section 3: Conclusions

Summary

CA Risk Authentication enhances simple password mechanisms by evaluating the risk associated with the login in context to the individual user. This approach is transparent to the end user as it requires no action on their part while providing greater assurance that the user is who they claim to be. The majority of logins will be performed by legitimate users, and these will continue uninterrupted because their risk score is low. And for the few higher-risk scenarios, CA Risk Authentication provides a simple and easy means to validate a user's identity. In addition, user behavioral profiling ensures that risk is evaluated based on what is normal for each individual user and can detect when their behavior deviates from the norm. Finally, risk can be assessed throughout the user's session, especially when the user is about to conduct a sensitive transaction. This guards against stolen cookie attacks.

CS200-171227_1215