RSA DISTRIBUTED CREDENTIAL PROTECTION

There is a security weakness lurking in many of today's best designed systems – a primary point of compromise. Think about your own IT operations. Chances are that by breaching just one critical authentication server, application, or data store an attacker can gain access to your most sensitive data.

RSA® Distributed Credential Protection (DCP) is designed to help eliminate primary points of compromise by scrambling, randomizing and splitting secrets and authentication decisions across two servers. Even if one server is breached, DCP is designed to help ensure that both authentication decisions and data remain secure.

White Paper

THE HAZARDS OF PRIMARY POINTS OF COMPROMISE

Most systems today—even those designed using best security practices—have a security weakness - a primary point of compromise. An attacker only has to compromise one server, piece of software, or machine to access the system's most sensitive data. If an organization has sensitive data in a system with a primary point of compromise, it bears all of the potential liability resulting from an internal breach.

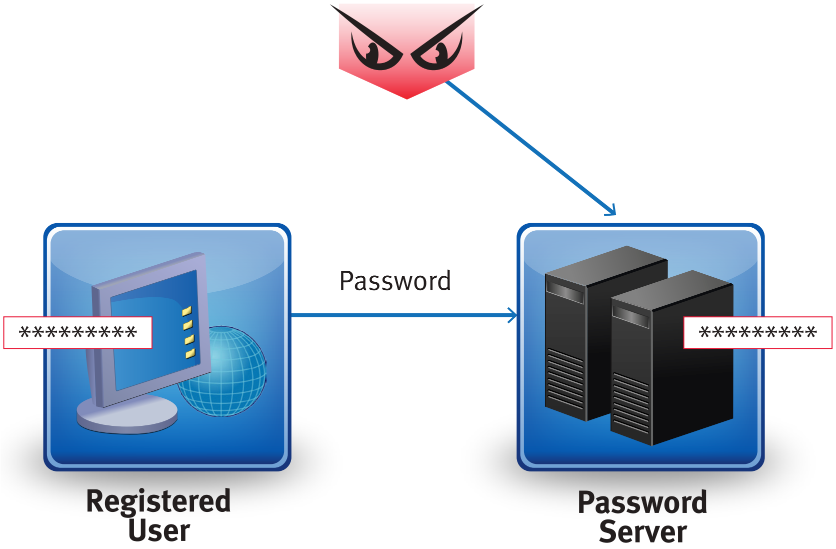

Figure 1: Primary Single Point of Compromise

Figure 1 illustrates a primary point of compromise that exists today. A typical portal or application requires a user name and password in order to register and then subsequently log in. Today, that password is most likely stored in a central password server or authentication server.

This is a primary point of compromise for an attacker. All of an enterprise's credentials or passwords can be found here; even if the credentials are hashed and/or salted, attackers have proven that there are ways around this in order to reveal the plaintext information.

HOW RSA DISTRIBUTED CREDENTIAL PROTECTION WORKS

RSA Distributed Credential Protection's basic premise is simple. It includes two servers (Server 1 = Blue and Server 2 = Red), that jointly verify authentication attempts. Credentials such as passwords are cryptographically scrambled, randomized and split across the two servers. If one server (e.g., Blue) is corrupted by an attacker, the other server (e.g., Red) continues to help ensure the integrity of the authentication process. The credential data in Blue or Red by itself reveals no information about the credentials themselves.

RSA Distributed Credential Protection can protect authentication through passwords, knowledge-based questions (e.g., "What's the name of your first pet?"), or certificates. Together, Blue and Red form what's called the authentication plane in Distributed Credential Protection.

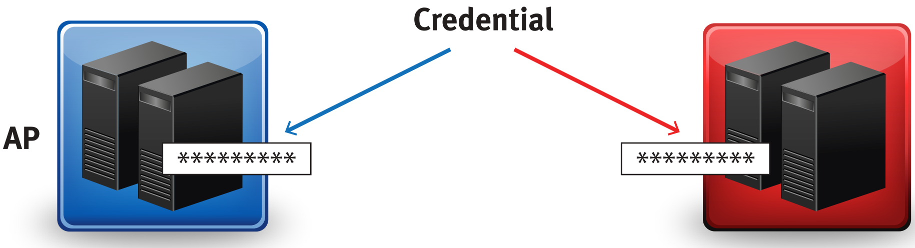

Figure 2: Authentication plane

Figure 3 depicts the process by which a user, Alice, gains access to an application using a password protected by DCP. She sends her password (randomized and in split form) to Blue and Red in the authentication plane. Blue and Red jointly verify this password. If it's correct, they allow Alice to access the application. If it isn't correct, access is denied.

Figure 3: Gaining access to secrets with RSA Distributed Credential Protection

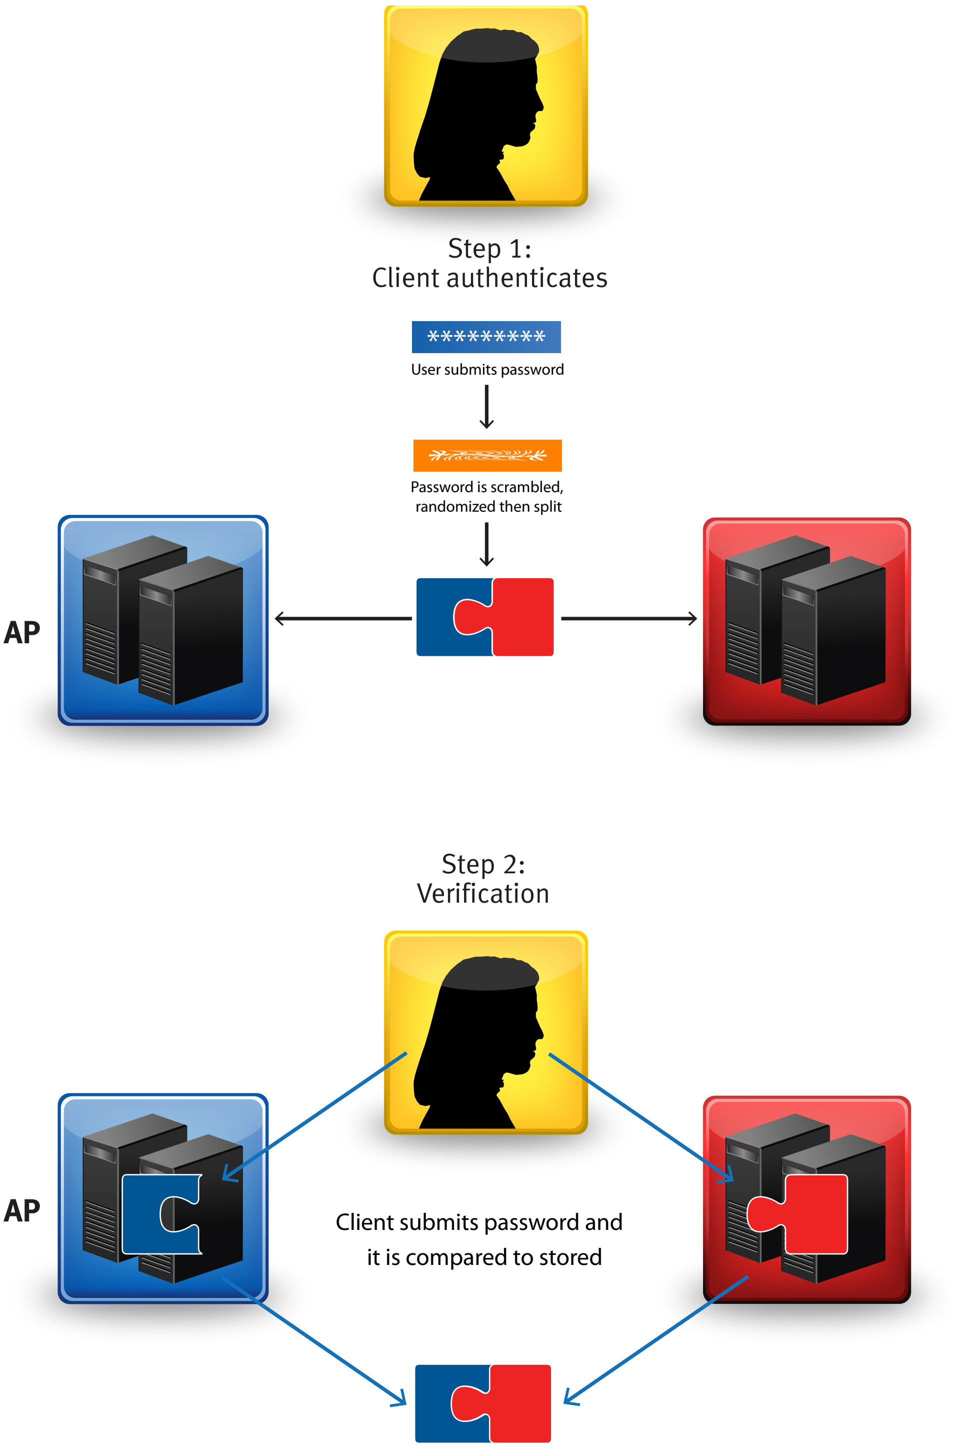
Step 1:
Client authenticates
****
User submits password
Password is scrambled, randomized then split
AP
Step 2:
Verification
AP
Client submits password and it is compared to stored

As stated previously, DCP also allows for re-randomization of the secrets. RSA Distributed Credential Protection periodically refreshes the randomness used to split credentials. This approach, known as proactive cryptography, enables recovery from transient breaches, even those that go undetected. An attacker that breaks into Blue and only breaks into Red at a significantly later time achieves no more than by breaking into Red alone.

Distributed Credential Protection applies proactive cryptography by computing a new random pad T and applying T to the values stored on both Red and Blue. On Blue, Alice's initial password P, which has been masked by R, is now also masked by T. The value P+R+T is stored on Blue, replacing P+R. Over on Red, the original pad R is replaced by R+T. As before, both P+R+T and R+T are, on their own, completely random. (Also as before, if they are added together, they are equal to P.)

This re-randomization can occur on an automated schedule, or on-demand in the event of a breach.

Figure 4 illustrates this re-randomization.

Figure 4: Re-randomizing protected secrets

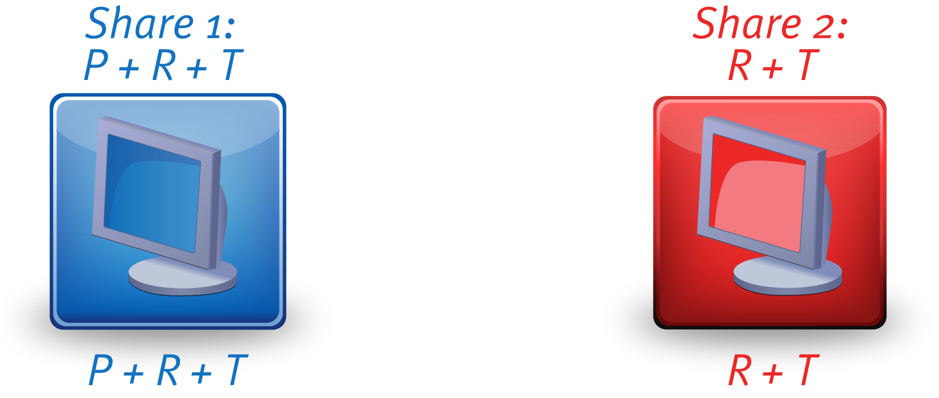

THE MAGIC OF CRYPTOGRAPHY: VERIFYING SECRETS WITHOUT SEEING THEM

The authentication plane in DCP can work with static secrets such as passwords or answers to knowledge-based questions. But when a user, Alice, authenticates in DCP using a secret P, neither Blue nor Red ever handles P as cleartext. This is important for two reasons. First, it's an essential property in protecting the authentication plane against breaches. Additionally, though, it helps allow for secure use of PII for authentication: in DCP, such PII is never exposed.

How do Blue and Red authenticate passwords without reconstructing them? Here's how the cryptographic magic works.

With RSA Distributed Credential Protection, passwords are stored in a randomly split form. When user Alice registers her password P, the DCP agent also computes a random pad R. The password masked with the pad, P + R, is stored on Blue, while the pad R is stored on Red. Either value by itself is completely random. (Added together, though, they are equal to P.)

When Alice authenticates to the system by entering a password P2, the DCP agent splits the password again with a fresh, random pad S. The value P2 + S is sent to Blue and S is sent to Red.

In order to authenticate, Blue computes A = P + R – (P2 + S) = (P – P2) + R – S, while Red computes B = R – S. If Alice has entered the correct password, then P = P2 and A = B; that is, Blue and Red hold the same value.

To verify passwords, then, Blue and Red simply need to check whether or not they're holding equal values. Cryptographic techniques developed by RSA Laboratories enable them to perform blind equality testing of this kind securely and efficiently.

Figure 5 depicts the process of authentication in DCP for a personal question of the type often used for password recovery ("What is the name of your first pet?").

Figure 5: Authentication process with RSA Distributed Credential Protection

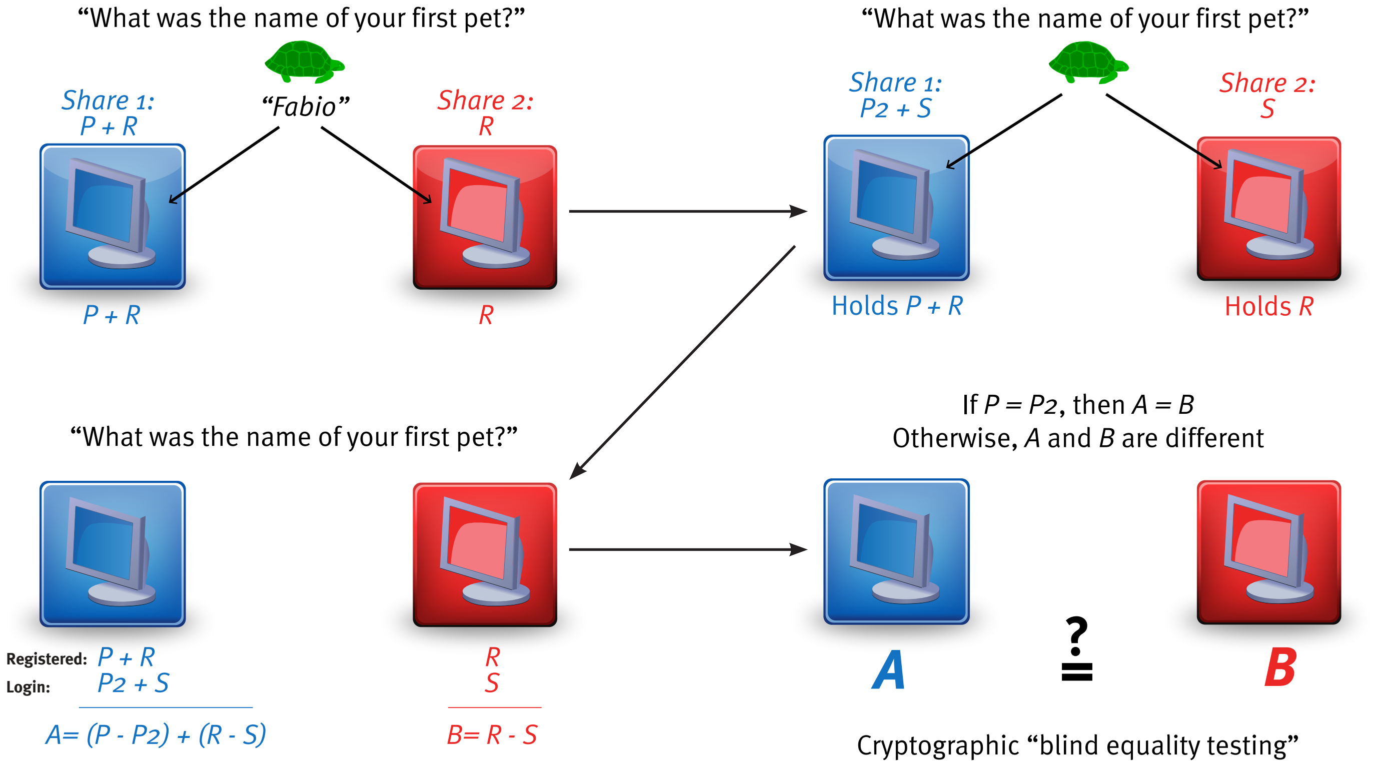
"What was the name of your first pet?"
"Fabio"
Share 1:
P + R
Share 2:
R
P + R
R
"What was the name of your first pet?"
Share 1:
P2 + S
Share 2:
S
Holds P + R
Holds R
"What was the name of your first pet?"
Registered: P + R
Login: P2 + S
A= (P - P2) + (R - S)
R
S
B= R - S
If P = P2, then A = B
Otherwise, A and B are different
A
?
=
B
Cryptographic "blind equality testing"

FLEXIBLE DEPLOYMENTS AND MIGRATION PATHS

To ensure the strongest breach resistance, the Blue and Red server configurations can be diversified. This diversification can take a variety of forms, offering different security and ease-of-deployment tradeoffs.

In its most basic form, RSA Distributed Credential Protection can be deployed on a single physical server with Blue and Red sitting in distinct virtual machines (VMs). Use of different operating systems across VMs prevents an attacker from exploiting a single OS vulnerability. This option offers an easy stepping stone to more sophisticated DCP deployments.

Stronger security in DCP comes with greater diversification. Blue and Red can reside in separate administrative domains or even in separate organizations. Particularly attractive is an enterprise, operating Blue, cooperating with a cloud provider that operates Red. The cloud provider then helps the enterprise prevent data loss or authentication failures in the wake of an internal breach. Figure 6 illustrates the spectrum of deployment options for DCP.

Figure 6: Spectrum of deployment options

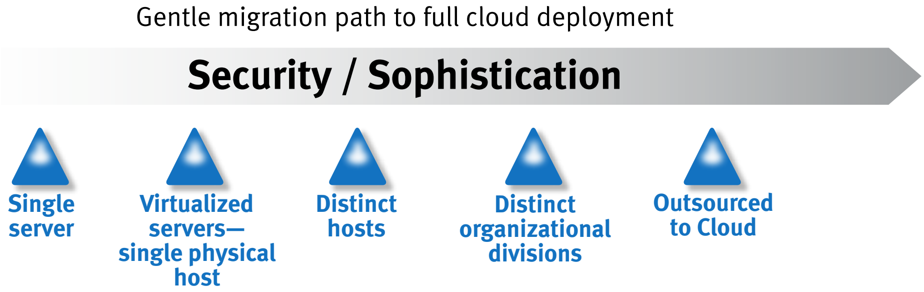
Gentle migration path to full cloud deployment
Security / Sophistication
Single server
Virtualized servers— single physical host
Distinct hosts
Distinct organizational divisions
Outsourced to Cloud

WHAT RSA DISTRIBUTED CREDENTIAL PROTECTION CAN DO FOR YOU...

Regardless of what kind of data you hold or how users authenticate to your network, DCP can bolster your defense-in-depth strategy with a strong new layer.

By isolating or even nullifying the impact of a breach, DCP can greatly help reduce this risk for your organization. If an attacker does manage to penetrate one of your systems, DCP can automatically or manually refresh protections on your secrets and credentials to limit the value of exfiltrated data. DCP helps protect against detected breaches, silent system compromise, and malicious insiders alike.

CONCLUSION

Server compromise has almost become a de facto fact of life. Intruders are always going to attempt to attack servers. If you can't stop this, then you need something that can remove the impact of server compromise. By randomizing and splitting secrets and authentication decisions across two servers, and splitting data in a recommended data plane, RSA Distributed Credential Protection makes this a recoverable and manageable event instead of a potential news headline.