Skyhigh Security

WHITE PAPER

The Promise and Reality of Remote Browser Isolation

Browser protection without detection

Table of Contents

The modern paradigm of ubiquitous interconnectivity has transformed web traffic into one of the most popular delivery vehicles for today's sophisticated cyberattacks. The high volume and complexity of these attacks has challenged the security industry to find more dynamic and intelligent methods to protect organizations against them. One of the more recent solutions from the security industry is Remote Browser Isolation (RBI). This new technology takes an entirely different approach, promising to ensure security through isolation rather than detection. RBI is becoming a common component of many web security offerings in the market. Many RBI technology startups have been acquired and integrated into mature Secure Web Gateway (SWG) solutions as an additional layer of protection. However, while RBI may be one of the industry's most exciting new technologies, a few critical factors separate expensive, unmet potential from a true paradigm shift in webbased security.

What is Remote Browser Isolation?

Before looking at the benefits and challenges of RBI, let's look at the technology and the most common use cases. When a user requests a website, RBI will render the requested content in a temporary browser in a remote data center and then allow the user to view and interact with the content. This is usually transparent to the user, providing a typical browsing experience but without loading the requested content on the user's local machine. There are various implementations of this technology, but the core principle is that potentially unsafe content never reaches the user's endpoint. The fundamental value of this approach is that protection is no longer contingent on detection. Malware may go undetected, but it will only compromise the temporary browser that has no access to valuable assets. Therefore, endpoints are protected regardless of successful detection.

While the benefits of this technology are apparent and the need to protect corporate endpoints is universal, there are two ways in which organizations use RBI. By far, the most common model is to protect against uncategorized sites or sites that have unknown risk. All web protection vendors will have gaps in their intelligence—even state-of-the-art threat protection stacks can be plagued by hundreds of potentially dangerous "false negatives" every day. As a result, organizations are often caught in a catch-22 when deciding how to handle these "gray areas." If they block all uncategorized sites, organizations must painstakingly maintain an allow list of necessary sites that are not known to be safe. Alternatively, if they allow these sites, this presents a tremendous risk of a newly registered site being malicious and delivering a web-based attack. RBI is a "best of both worlds" solution to this problem.

In addition, many organizations have a group of users that need an extra level of protection at all times. A myriad of scenarios may require this, such as users who have access to highly sensitive data or users who have elevated privileges. For example, many organizations choose to isolate all traffic for C-level executives and IT administrators. Some users, such as malware researchers or incident response personnel in the security operations center (SOC) may knowingly visit potentially dangerous sites as part of their research. In all these cases, users may require full-time isolation of web traffic, rather than just selectively isolating risky or unknown websites.

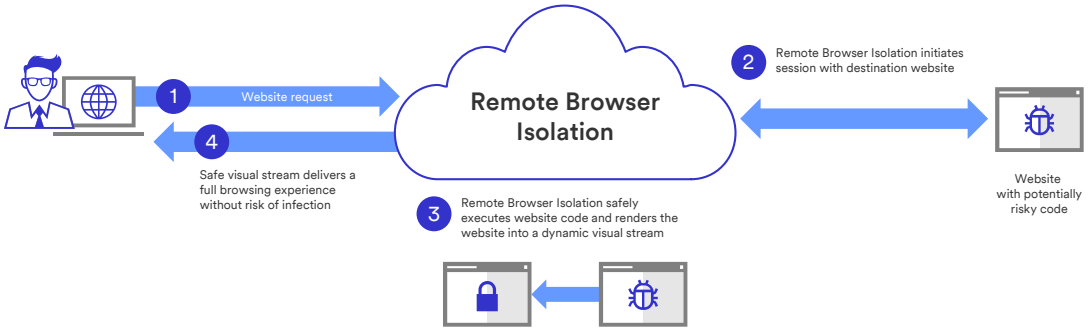

Figure 1. An RBI session in action.

Does the Need Worth Justify the Cost?

One common characteristic of all full-time RBI solutions is the expense. RBI is costly for vendors to deliver because rendering web content requires tremendous infrastructure resources to do at scale. An instance of Chrome with a single tab open can often consume nearly a gigabyte of memory. Consider how many tabs a single user may have open on average, and you begin to understand why the technology doesn't come cheap. In some cases, the cost of deploying an organization wide RBI solution may consume or even exceed an organization's entire security budget.

In light of that, consider the most common use case of selectively isolating only uncategorized sites. Some vendors require licensing all users for full-time isolation, without a less expensive option for selective isolation of the "gray areas." This means an organization will be forced to pay the full cost of RBI when it will likely be used for as little as 1% of their web traffic. While this is a powerful tool to address a critical need, few CISOs will be able to justify spending such a large portion of their security budget on such a small percentage of their web traffic.

In answer to this, some RBI vendors have begun to offer licensing for selective browser isolation of risky or unknown websites, but the discount will almost certainly fail to be proportional. For example, selective isolation licenses might be offered at a 60% discount, but that still fails to align with the fact that 1% of those users' web traffic is being isolated.

To better understand these cost implications, consider a hypothetical use case of an organization with 10,000 users. Let's assume that IT has designated 400 "priority" users that will require web isolation for all of their traffic, while the rest would only need selective isolation applied to roughly 1% of their web traffic that is uncategorized.

In this case, while high-risk users only comprise 4% of the overall user population, they represent over 80% of the traffic being processed by the RBI solution. At a per-user cost of $100, even assuming that the licensing for selective isolation is discounted to $40 per user, the overall solution cost would come out to $424,000, or a whopping $42 per user. Although they are responsible for less than 20% of all isolated traffic, the selective isolation users represent over 90% of the cost.

This excessive cost is incurred when pricing of enterprise-wide RBI deployments is based on number of users rather than solution utilization. An ideal RBI deployment that meets the requirements for most organizations will actually be a blend of full-time and selective isolation. Limited security budgets will dictate that any vendor selected must provide multiple licensing options, allowing organizations to license users according to their security requirements. These options must provide pricing that correlates with the security value realized. Selecting such an offering will enable organizations to strike the right balance between the cost of the solution and the security value realized from its use.

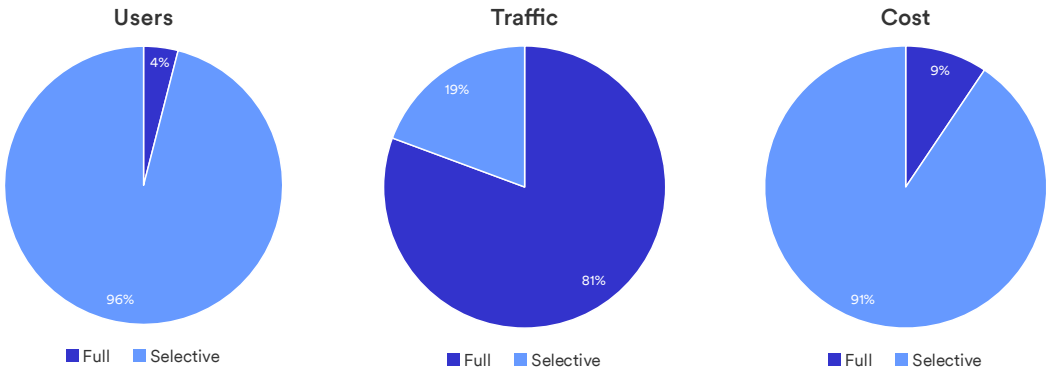

Figure 2. The disproportionate relationship between RBI users, traffic load, and solution cost.

Secure Web Gateway Convergence

Most RBI solutions on the market rely on integration with an SWG solution to help determine whether to isolate certain sites based on risk. This integration typically takes place using a simple HTTP redirect or "proxy chaining." Even if both your SWG and RBI solutions come from the same vendor, they may behave as two separate solutions with this type of loose integration.

The details of these integrations can introduce challenges with authentication, reporting accuracy, and a transparent user experience. When simply using an HTTP redirect, the SWG solution no longer has visibility into the web traffic being generated. This will "blind" the SWG, resulting in reporting gaps that can only be filled by merging log data from RBI. This is often true when using proxy chaining as well because RBI traffic frequently uses a proprietary protocol or WebSockets, making the traffic between the client and the remote browser impossible to decipher. When using an HTTP redirect, this is often very conspicuous to the user and may even require manual authentication to proceed.

From a deployment perspective, there may be misalignment between the SWG and RBI solutions. For example, the organization may insist that the solution has a robust, global cloud footprint; adhere to a service level agreement; or maintain certain security certifications, such as ISO 27001 or SOC2. Many customers may find that one solution meets their deployment objectives while the other does not.

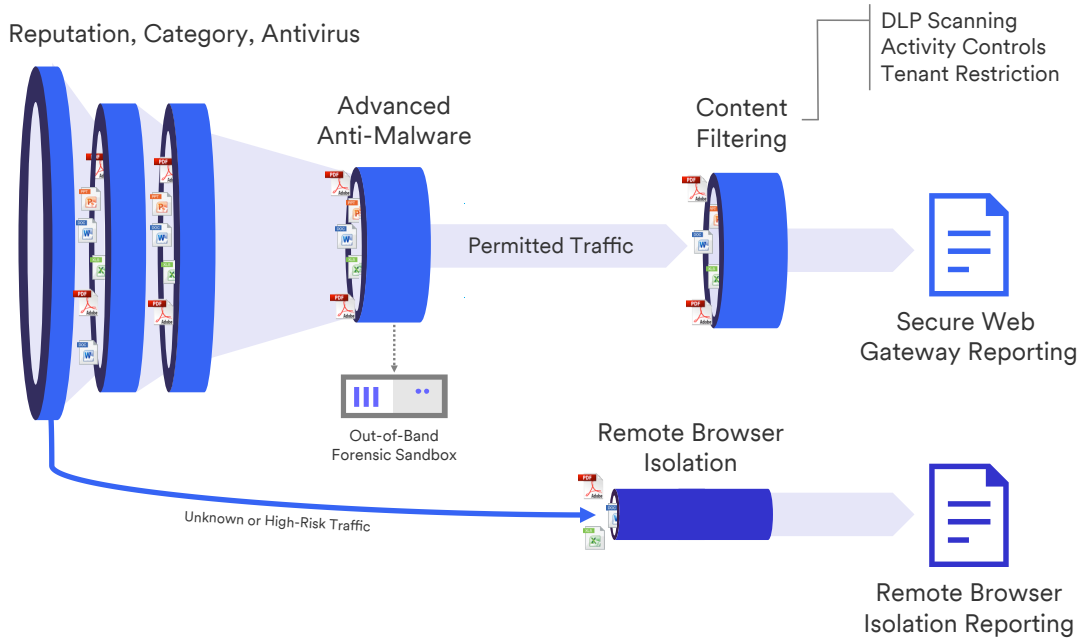

Figure 3. Loose integration between SWGs and RBIs.

Another undesirable outcome is management of two policies. Two solutions with such a loose integration are unlikely to share policies or even feature consistent filtering capabilities. For example, URL categories may not align properly between the policies, or DLP classifications may be defined differently. This will result in management overhead for already overburdened security teams and can create a confusing and inconsistent experience for the user.

Selecting a solution featuring RBI fully converged with a mature Secure Web Gateway will provide a single security platform that leverages a single unified policy, consistent data protection between isolated and nonisolated traffic, and a seamless administrative and user experience. This eliminates many of these potential challenges and yields a streamlined, "best of both worlds" outcome.

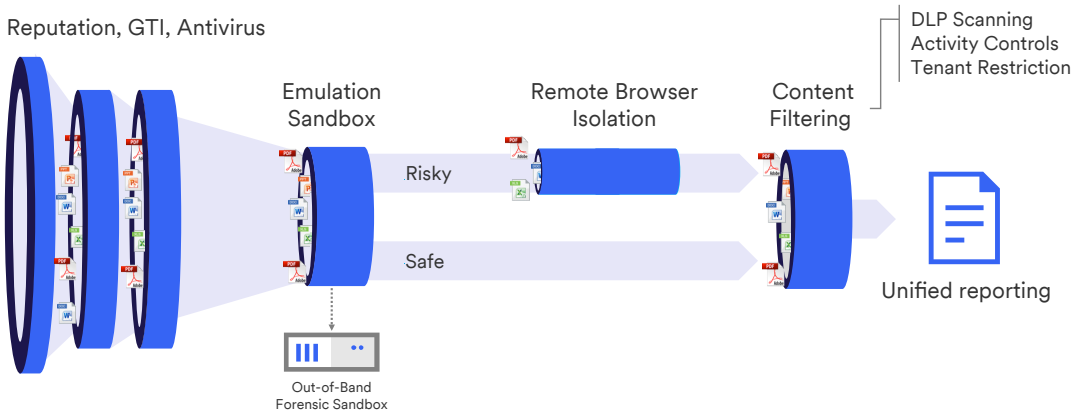

Figure 4. Integrated RBI with in-line data protection and activity controls.

The Human Factor

The SWG-RBI convergence is equally valuable when considering risks stemming from the human factor. The user behind the screen is still entirely susceptible to being duped by social engineering attacks even when delivered through this isolated medium. Many of the controls typically used to protect users might actually be forfeited when choosing isolation. SWGs are often wellsuited to protect organizations and users from social engineering attacks by preventing risky interactions with untrusted content. For example, an uncategorized website might be allowed, but login attempts by the user may be blocked by preventing any HTTP POST (which requests the web server to accept the data enclosed in the body of the POST message) to the site. This is an effective protection against fake login pages that may be phishing for users' credentials.

Additionally, a lack of sensitive data awareness may be the death knell for RBI if the goal is full isolation for certain users. When selectively isolating traffic to uncategorized sites, it's common practice to block all uploads, in which case sensitive data awareness is not required. However, when isolating all traffic as earlier described in the full isolation use case, sensitive data awareness is paramount to prevent risky uploads to otherwise trusted sites. For example, universally blocking uploads to a sanctioned application is untenable, but allowing uploads while scanning for proprietary or regulatory data allows use of RBI while maintaining data security visibility and control.

In order to compensate for the human factor, it's vital to look for an RBI solution that features in-line data protection and activity controls. Look for these capabilities either built into the RBI solution itself or through a tight integration with a SWG security stack as shown below.

Protection Without Detection

At its very essence, RBI provides protection without detection. By rendering web-based content as a dynamic visual stream on a remote browser that users can safely interact with, malware risk is transferred to an asset that has effectively no value. This completely insulates and protects an organization's assets from any webbased malware without any reliance on the detection of potential threats. That's the true value of the technology.

There are two competing varieties of browser isolation on the market: pixel-mapping and document object model (DOM) mirroring. Pixel-mapping loads a full, remote browser and renders all web content in that browser. This method fully and absolutely isolates endpoints from all content from the requested site. DOM mirroring selectively renders some parts of the requested content remotely, while other parts are scanned and then allowed through in their original format. Vendors that use this approach claim that it improves performance, but a more likely agenda is to reduce cost by only partially rendering content remotely. Security teams should be wary of any RBI technology that claims to "dissect" or "scan" web content and then send the safe content through while isolating the risky content. These terms point to the approach of using detection to distinguish the safe content from the risky content. At its very core, this violates the inherent value of browser isolation, which is protection without detection.

Another argument commonly used to support the DOM mirroring method is that by dissecting web content and selectively rendering content, greater visibility and context is achieved by the RBI solution enriching logging and reporting. This is a valid point, but it's rendered moot by true convergence with an SWG solution. As outlined earlier, a properly converged solution, where RBI can be leveraged without losing proxy-level visibility, ensures that a pixel-mapping RBI implementation will maintain similar or better visibility and context when compared to DOM mirroring.

Introducing Skyhigh Security Remote Browser Isolation

No RBI solution search would be complete without considering Skyhigh Security. Our Remote Browser Isolation is the first solution in the industry to be seamlessly converged with SWG, Cloud Access Security Broker (CASB), Data Loss Prevention (DLP), and Zero Trust Network Access through our Security Service Edge (SSE) solution.

Key advantages:

- Comprehensive DLP: Enhanced visibility and protection over how data is being accessed or shared.
- Part of a complete threat protection stack: Works directly in-line with our SSE threat protection, ensuring consistent policies, data protection, and visibility across isolated and non-isolated traffic.
- Simple to use: Works seamlessly with standard web browsers, so users require no training or changes in behavior.
- Fast and responsive: Websites load quickly and are immediately responsive to typing, clicking, and scrolling— no more slow web browsing.
- Powerful management: Robust policy and reporting engines provide the optimal flexibility and granularity to secure users' browsing activities.

Skyhigh Security Remote Browser Isolation for risky web traffic is provided for no extra charge as part of our SSE solution, with the opportunity to incrementally add full isolation licenses for those users that need it. This completely disrupts existing cost models for enterprisewide RBI deployments and makes the next great security technology attainable for any organization.

Conclusion

Remote Browser Isolation is perhaps the most innovative weapon developed to fight web-based malware to date. While there are several important considerations with respect to this new technology, they should not discourage exploration. Rather, these considerations are key criteria that need to be assessed when comparing the numerous solutions on the market. RBI still promises nearly unbeatable protection against web threats, which certainly justifies some diligence in finding a solution that's ideal for any organization.

About Skyhigh Security

When your sensitive data spans the web, cloud applications, and infrastructure, it's time to rethink your approach to security. Imagine an integrated Security Service Edge solution that controls how data is used, shared, and created, no matter the source. Skyhigh Security empowers organizations to share data in the cloud with anyone, anywhere, from any device without worry. Discover Skyhigh Security, the industry-leading, data-aware cloud security platform.

For more information visit us at skyhighsecurity.com

Skyhigh Security

3099 North First Street
San Jose, CA 95134
888.847.8766
skyhighsecurity.com

November 2023. WP-EN-000013-01