Remote Incident Manager

Executive Summary

The Remote Incident Manager (RIM) has provided secure, efficient, accessible, feature-rich remote desktop access, enabling field technical support providers to cut costs while improving quality of service since 2007. Now we bring the power of Pneuma Solutions accessible cloud platform to the RIM community. This is a detailed but not thorough explanation of RIM. Ask your Pneuma Solutions representative any questions addressing your specific use case not addressed here.

Introduction

The Power of Remote Desktop Access

Remote desktop access enables a user to take complete control of a distant computer over an organization's network or the Internet, as if the user was sitting at that computer. This technology has a wide range of uses; some of the most common uses are:

- technical support and troubleshooting
- training and distance learning
- system administration

When combined with real-time voice communication, remote desktop access makes location irrelevant for most support and training tasks. Technicians need not face the frustrations of explaining complex step-by-step procedures to end-users; instead, a technician connects to the user's computer and interacts with it directly to solve the problem. Travel expenses are reduced, and technical problems are resolved more quickly. Field technical support providers need not maintain a full staff in each location they serve. A competent individual or small team can perform most field technical support tasks for a large, scattered user base. In short, remote desktop access enables field technical support providers to cut costs while improving quality of service.

The Unrealized Potential

Though remote desktop access is a powerful tool for technical support, the current remote technical support products have serious limitations that prevent this technology from reaching its full potential.

In particular, the most widely used remote technical support products, such as TeamViewer, don't enable accessibility for blind technicians who must use a screen reader. Even though these products generally mirror the audio being played on the remote computer, they still require a screen reader to be run on that computer, meaning that the end-user will hear the speech output, even if the end-user is sighted and doesn't need or want speech output. This prevents skilled blind technology professionals from working as

remote technical support providers without having to disclose that they are visually impaired to the end user.

Mainstream remote technical support products are also inadequate for supporting blind end-users who are already running a screen reader. Even though these products mirror the remote audio, there are some issues that make remote desktop products incompatible with some screen readers running on the remote machine. For example, when trying to use JAWS on an end-user's machine in conjunction with one of these popular remote desktop products, the remote instance of JAWS ignores keyboard commands from the technician, unless JAWS is reconfigured on the end-user's machine. This makes it very difficult for blind technicians to support blind end-users running JAWS, but also makes it harder for sighted technicians to support JAWS users.

While JAWS supports Microsoft Remote Desktop through an optional paid add-on, this feature isn't usable for most remote access scenarios. Microsoft Remote Desktop isn't designed to let a technician connect to the existing console session on a remote machine. Instead, a Microsoft Remote Desktop connection starts a new session, with its own screen and running applications, and forces the user sitting in front of the remote computer to log out. Also, even with the JAWS remote desktop add-on, JAWS still has to be installed on all machines that the technician might need to remotely access. So this solution may only be practical for remote access to a small number of servers, not desktop PCs.

Remote access solutions that are integrated with third-party screen readers, such as JAWS Tandem and NVDA Remote, can only be used if the end-user is running the same screen reader as the technician. NVDA Remote doesn't even mirror the remote user's screen, making it unusable for sighted and low-vision technicians.

These screen-reader-specific solutions also fail to take advantage of the latest peer-to-peer networking technology, with cloud-based relays as a fallback, to allow the technician to connect to the end-user's machine in a way that's both secure and responsive without requiring cumbersome and invasive changes to network configurations. Both JAWS Tandem and NVDA Remote provide two options: a central relay service, or direct connectivity that requires a port to be opened in the firewall on the end-user's machine and network. While the central relay service is convenient, it increases the latency of the remote session, particularly when the technician and end-user are far away from the relay. The relays for both JAWS Tandem and NVDA Remote are deployed in a single geographic region. While NVDA Remote's relay server implementation is open-source, and third parties can operate their own relay services, this requires both the technician and end-user to configure NVDA Remote with the address of the alternate relay service.

An Inclusive Solution: Remote Incident Manager

The Remote Incident Manager (RIM) is an innovative remote desktop access package designed for field technical support, which resolves all of the problems described above, while providing all of the features that today's IT professionals expect from remote technical support products.

Deployment, Management, and Security

At the center of RIM is a cloud service which manages all remote access connections. On-premises and private cloud deployments are also available for enterprise customers. Network security policies need not be adjusted for each remotely accessible machine. The RIM service provides a Web-based interface for centralized deployment and management of remotely accessible machines (called targets). At the same time, connections between technicians' machines (called controllers) and targets are peer-to-peer when possible, and always encrypted end-to-end. In corporate environments of all sizes, this means that remote connections can be securely made across both the corporate network and the Internet, whether the customer is using the RIM public cloud service, or an on-premises or private cloud deployment.

For enterprise customers, targets can be divided into target groups, which are registered on the RIM central service (described below). A target group may correspond to a department, office, or other organizational unit. Target groups are used to organize large numbers of machines, to set policy, and to ease deployment of the RIM Enterprise app on target machines.

In short, the RIM service offers the convenience of centralized deployment and management without sacrificing security or accessibility.

Shared Session

Like all products designed for remote technical support, RIM puts the technician in the same Windows session as the end-user. Not only can the two people work with the same applications and documents at the same time, they can also exchange text and even files through the Windows clipboard. This is invaluable for field technical support because the technician can watch the user and show how to perform required tasks or avoid common problems.

Proactive Field Technical Support

While RIM makes it easy for an end-user to set up their machine for remote access on the spot, a system builder, computer technician, or assistive technology specialist can also choose to install RIM on the end-users' machines in advance, so users can easily activate the feature at any time, even in many circumstances that render the computer otherwise unusable. Users may also opt to allow a Controller access to their machine at a designated time to perform routine maintinence and/or software upgrades without requiring the user to be present at the machine. When a user has a problem, the technician need not specify the desired remote machine by IP address or host name; instead, the technician and user both enter a session keyword chosen by the technician. If the technician needs to restart the user's computer while solving the problem, they can choose to automatically re-connect to the computer once it is restarted; the technician can even force the computer to restart if it is unable to carry out the normal shutdown procedure. Thus, RIM makes it easy for a technician to serve remote users.

The option to proactively set up RIM for easy remote technical support makes it an especially good option for organizations, such as vocational rehabilitation agencies and

training centers, that provide pre-configured computers to their end-users. Such organizations can easily install RIM on all end-user computers in advance, so it's ready to go whether or not the organization itself provides remote technical support. This allows the user to take advantage of a number of blind IT professionals that are familiar with assistive technologies.

Accessibility

RIM is fully accessible to all technical support professionals, whether they're blind, low-vision, or sighted. RIM is affordable regardless of user base size, with flexible pricing for both individual technicians and organizations of all sizes. Sighted end-users are given no direct indication that the technician is blind or low-vision; they do not hear the speech output that the technician requires. There is no need to install a screen reader on end-user machines in advance, or to make any visible change to the end-user's Windows configuration. The technician receives speech and/or Braille output through an included custom add-on for the NVDA screen reader.

In short, RIM is the most accessible remote desktop access package on the market. Field technical support providers can now easily and cost-effectively benefit from a large and growing pool of competent professionals who were previously unable to use remote desktop access products.

Responsive Remote Access

RIM takes advantage of the latest peer-to-peer networking technology to give technicians highly responsive remote access to end-user machines. Whenever possible, RIM will establish a direct connection between the two machines, without requiring any cumbersome changes to firewalls or other network configuration details. When a direct connection is not possible, RIM will intelligently use our cloud relay service. Because we deploy relays in several locations around the world, the remote session is responsive even when a peer-to-peer connection is not possible. For enterprise deployments with highly restricted Internet access, RIM can also be configured to use one or more relay servers on the corporate network. And as noted above, the connection is always encrypted end-to-end, so Pneuma Solutions can't snoop on the session even when a cloud relay is being used.

High-Quality Audio

RIM lets the technician hear all audio playing on the end-user's machine. Using the best available real-time audio encoding technology, this audio is sent across the connection in full-fidelity stereo. This feature is useful not only for supporting blind users running screen readers, but also for supporting end-users in audio-based applications such as music production.

Integrated Voice Communication

The technician may choose to start a real-time voice conversation with the end-user as part of the remote session. This feature builds on the same robust WebRTC technology used by many meeting and chat apps to provide reliable, low-latency voice communication,

complete with echo cancellation, so neither the technician nor the end-user need to wear headphones.

RIM Components

A RIM remote session consists of three components: the controller, the target, and the central service.

Controller

The RIM app running on the technician's machine is called the controller. Technicians can quickly and easily gain access to the RIM service on any Windows computer they encounter by simply installing the app from the RIM website, or from an on-premises or private cloud deployment of the RIM service for enterprise customers. If the technician is using the app on a machine that's not their own, then it will automatically remove itself from the machine at the end of the session.

Target

The same app is installed on end-user (target) machines. The installer gives system builders and technicians flexibility in deployment options. Once deployed in the target role, the app sits quietly in the background, consuming minimal machine resources except when remote access is required. It makes no permanent system configuration changes, except to install itself as a Windows service. Except for a small icon in the system tray (normally located at the bottom of the screen), users will not notice that the app is present until they need it.

The target machine plays a crucial role in RIM's accessibility. Because the RIM app sends all audio being played on the target machine, the technician can hear the speech output from any screen reader the end-user is already running. If the end-user is running JAWS, Narrator, NVDA, SuperNova, or System Access, the technician will be notified about this at the start of the connection. If the end-user is not running a screen reader, then the target machine will coordinate with the controller to provide speech and/or Braille output to the technician, while not requiring the end-user to hear speech if used by the technition.

Central Service

The central service is either the RIM public cloud service for most customers, or an on-premises or private cloud deployment for enterprise customers. In either case, this service manages all controllers, targets, and connections between the two. It provides an easy-to-use, fully accessible, Web-based interface for all management tasks. The service can also email installation instructions to target users on the technician's behalf.

The central service plays a vital role in RIM's security. All connections between controllers and targets are coordinated by the central service. This means that no one can connect to a RIM target machine except through the central service.

Once the connection is established, RIM is designed to minimize the ongoing involvement of the central service, for maximum reliability and responsiveness. As explained above, RIM

connections are peer-to-peer whenever the controller and target networks allow peer-to-peer connectivity. When a peer-to-peer connection isn't possible, RIM uses our worldwide cloud relay network. For added reliability, this network is independent from the central RIM cloud service. Enterprise customers can also choose to deploy their own relay(s). In all cases, the connection is encrypted end-to-end, so the central service will never compromise security and privacy.

As mentioned earlier, an enterprise customer can either install and run its own private server or deploy on a corporate private cloud. Deploying a private server is straightforward, and a private server provides the greatest control over security. Depending on the size of the network, running a private server may be more cost-effective in the long term than using the RIM public cloud service.

Use Cases

Controller Setup (Public Cloud)

1. The technician downloads the installer from the RIM website.
2. After downloading the installer, the technician runs it.
3. A new installation of RIM is initially configured for the target role, so it prompts for a session keyword. However, the same screen has a button called "Provide Help Instead", which the technician presses at this point.
4. RIM then asks the technician whether they are installing it on their own PC. If they are not, RIM will automatically remove itself from the PC after the session has ended.
5. RIM prompts the technician to log into the cloud service using their email address.
6. The next step depends on how the technician or their organization has chosen to handle authentication with RIM.
    - If the technician's organization has integrated RIM with their single sign-on (SSO) service, RIM will then open the technician's default browser and redirect them to the registered SSO provider. When the SSO process is complete, the technician's browser is redirected back to the RIM cloud service, which provides a button that the technician can use to return to the RIM app.
    - Otherwise, the cloud service will send an authentication code, to the technician's email address or registered phone number, that they can use to complete the login process.
7. The technician is now ready to deploy targets or start a remote session.

Controller Setup (On-Premises or Private Cloud)

1. The technician logs in to the central service using any browser.
2. The technician downloads and runs the installer from the central service.

3. The RIM Enterprise desktop app is installed.

4. The app instructs the technician to return to the web browser where they downloaded the installer, and press a button to connect the app to the organization's RIM deployment.

5. When the user presses this button, the browser prompts them to confirm that they want to open a custom link in the RIM Enterprise app.

6. As soon as the user answers this browser prompt, RIM Enterprise is ready to use as a controller. The technician is now ready to deploy targets or start a remote session.

Target Deployment by the Technician

1. The technician runs the RIM app that they installed above.

2. On the main screen, they choose to deploy targets rather than starting a remote session.

3. For enterprise deployments, the technician may choose to create a new group for the targets being deployed, deploy to an existing group, or not use target groups.

4. The app provides a command line to pre-configure the deployed targets, which the technician can easily copy to the clipboard.

5. On the same screen, the app provides a button to go to the already downloaded installer in File Explorer. Note that RIM uses the same installer for both the controller and target roles.

6. Because the command line provided in step 3 enables the target to be installed and configured non-interactively, technicians have flexible deployment options. A technician serving a small user base may install the app on one computer at a time. A system builder may include the package in an automated system setup procedure. And a technician serving a large corporate network may use an automated deployment tool to push the installer and the provided command line to all computers at once. Refer to the next use case for information on installation by end-users.

7. Once installed, the app on each target machine automatically connects to the central service (either the RIM public cloud service, or an enterprise on-premises or private cloud deployment) and is ready for remote access. No additional action per target machine is required.

Target Installation by End-Users (Public Cloud Only)

1. The technician may either tell the user to download the app from the RIM website, or have the cloud service email installation instructions to the user.

2. The user downloads and runs the RIM installer.

3. The user is now ready to connect with the technician.

Field Technical Support or Training Session

1. The technician chooses a keyword for this session, such as the first name of the technician or user.

2. The technician starts the RIM app, enters the keyword, and sets other options for the session.

3. The first step for the end-user depends on how the app was deployed on the target machine.

    - If the app was deployed by the technician with a command line for pre-configuration, the end-user may start it by pressing Control+Shift+Backspace (or another hotkey configured by the technician) regardless of where they are on their PC, or by clicking an icon in the Windows system tray.

    - If the app was installed by the end-user, then it is started as soon as the installation is complete, and may be started later through the desktop or Start menu.

4. The app prompts the end-user for the session keyword. This prompt is both visual and easily accessible with a screen reader. If the target was deployed by the technician, they may configure the app to also provide a self-voicing prompt, to make it easy for blind end-users to start a session even if their normal assistive technology is not working.

5. The end-user enters the keyword provided by the technician. If a self-voicing prompt is configured as described above, the app echoes the characters audibly as well as visually.

6. Within a few seconds, the technician is connected to the end-user's computer and is ready to work.

7. RIM provides these options with regard to speech output:

    - If the technician needs speech output and the end-user's computer is not running a known screen reader, RIM automatically starts the accessibility component on the target machine, which then coordinates with the technician's NVDA installation to provide speech output. In this case, only the technician hears the speech output.

    - If a supported screen reader is running and functioning properly on the end-user's computer, both the end-user and the technician will hear that screen reader's speech output.

    - If a supported screen reader is running but not functioning properly, the technician can terminate that screen reader and start RIM's built-in accessibility component on the target machine. Again, only the technician will hear the speech output in this case.

8. If the technician needs to restart the target computer during a remote session, they will be automatically reconnected to the computer after it is restarted. If the computer is in such a state that the normal Windows restart process doesn't work, the technician can choose an emergency reboot option from the RIM menu to immediately reboot the target machine without performing a graceful shutdown.

9. Either the technician or the remote user can end the session. Both sides are notified when the session ends.

10. After disconnecting from the end-user's computer, the technician can leave a comment on the central service about the session, for record keeping and reporting.

Remote Access Session

1. The technician starts the RIM Enterprise controller on their computer.

2. The technician enters the host name or IP address of the remote machine and sets other options for the session. If the technician does not know the host name or IP address of the remote machine, they can ask the remote user to move to the RIM Enterprise icon in the system tray and read that icon's ToolTip.

3. The RIM Enterprise app on the target machine prompts the remote user to either accept or reject the technician's request for remote access. The target displays either the technician's name or a generic name such as "Help Desk"; the technician can configure this through the RIM Enterprise server.

4. The remote user presses Control+Shift+Y to accept the request or Control+Shift+N to reject it.

5. Assuming the remote user accepts the request, the technician is connected and ready to work within a few seconds. Both sides are notified when the connection is established.

6. Alternatively, through the RIM Enterprise deployment's web interface, the technician can configure a target group so that the end-user doesn't need to be prompted for access. This is ideal for managing servers through RIM Enterprise.

7. RIM Enterprise provides these options with regard to speech output:

    - If the technician needs speech output and the end-user's computer is not running a known screen reader, RIM Enterprise automatically starts its built-in accessibility component on the target machine, which then coordinates with the technician's NVDA installation to provide speech output. In this case, only the technician hears the speech output.

    - If a supported screen reader is running and functioning properly on the end-user's computer, both the end-user and the technician will hear that screen reader's speech output.

- If a supported screen reader is running but not functioning properly, the IT professional can terminate that screen reader and start RIM Enterprise's built-in accessibility component on the target machine. Again, only the technician will hear the speech output in this case.

8. If the technician needs to restart the target computer during a remote session, they will be automatically reconnected to the computer after it is restarted. If the computer is in such a state that the normal Windows restart process doesn't work, the technician can choose an emergency reboot option from the RIM Enterprise menu to immediately reboot the target machine without performing a graceful shutdown.

9. Either the technician or the remote user can end the session. Both sides are notified when the session ends.

Questions and Answers

Security

Do any ports need to be opened for the target machines?

No.

What types of incoming and outgoing connections must be allowed on the controller and target networks in order to establish a remote session?

On both sides, RIM must always first establish a connection to port 443 on the central service. For public cloud users, the hostname of this service is getrim.app. For on-premises and private cloud deployments, the hostname is decided by the organization at deployment time.

To establish the connection between the controller and target, RIM uses WebRTC, the same technology used by browser-based meeting platforms such as Google Meet, or the browser-based versions of Zoom and Microsoft Teams. The central service provides signaling for the WebRTC connection, but it does not relay actual session data. As is standard for WebRTC, RIM makes a peer-to-peer UDP connection whenever possible. When NAT traversal is needed, RIM uses the standard STUN and TURN protocols. In public cloud deployments, RIM accesses these protocols on our worldwide network of relays. For best results, outgoing UDP connections to ports 19302 and 3478 (the standard STUN and TURN ports) should be allowed; otherwise, RIM will fall back to TCP port 443. In enterprise deployments, RIM may be configured to use a relay or relays deployed on the corporate network.

Are remote sessions encrypted?

Yes; all remote sessions, including file transfers, are encrypted end-to-end using Datagram Transport Layer Security (DTLS), which uses the same technology as the HTTPS protocol used by all modern websites.

Can Pneuma Solutions eavesdrop on sessions relayed by the public cloud service?

No. Session key negotiation and encryption are performed end-to-end between the controller and the target; the cloud service merely relays data as-is. Therefore, the cloud service is unable to decipher the data that it relays. This also applies to file transfers.

What measures have been taken to prevent remote code execution (RCE) vulnerabilities?

All RIM code which is exposed to input from the network is written in memory-safe programming languages including Rust and JavaScript.

Does RIM comply with HIPAA?

Yes. For more information, please refer to our web site at www.PneumaSolutions.com.

Enterprise Deployments

Where can RIM Enterprise be deployed?

As an alternative to our public cloud services, enterprise customers may choose to deploy RIM on-premises or in a private cloud environment. We currently support virtual private cloud (VPC) deployments on Amazon Web Services and Microsoft Azure.

On which ports does the private server listen for incoming connections?

The private server listens on ports 80 and 443. Port 443 is handled using standard HTTPS and WebSockets, while the only purpose of port 80 is to redirect to HTTPS.

Does the private server require a server version of Windows?

No, the private server runs on all supported desktop and server versions of Windows and is also available for several common Linux distributions.

Does the private server require a database package such as Microsoft SQL Server?

No; the private server uses a built-in version of the SQLite high-performance, low-overhead, zero-configuration database engine.

Does the private server require a web server package such as Microsoft Internet Information Server?

No; the private server uses a built-in, high-performance, low-overhead web server.

Does the private server conflict with an existing web server on the same machine?

Yes; the private server listens on the standard HTTP and HTTPS ports. If this is an issue, the private server will need to be deployed on its own machine or VM.

Does the private server depend on any software apart from the operating system?

No; the private server is a self-contained package which will run on any Windows system with a no-hassle installation process, as well as on many common Linux distributions.

Does the private server require that its administrator have desktop access to the server machine?

No. The Windows package of the private server supports non-interactive installation, while the Linux package supports installation from the command line. After installation, all management is performed using a web browser.

What limitations exist on the number of target machines that can connect to the private server?

The private server imposes no hard limit on the number of target machines that can connect to it; this number is limited only by CPU speed, available memory, and bandwidth.

How does the private server handle TLS certificate setup for HTTPS?

The private server offers three options for setting up a TLS certificate:

1. The administrator can obtain a free certificate from Let's Encrypt. This requires the private server to have Internet access for both initial setup and recurring renewal.
2. The administrator can obtain a certificate from another service, perhaps internal, that implements the same ACME protocol as Let's Encrypt.
3. The administrator can import a custom certificate.

Conclusion

Remote desktop access is an immensely powerful tool for field technical support, regardless of user base size. The Remote Incident Manager addresses the problems that most hinder field technical support providers from harnessing the power of remote desktop access effectively. It provides security, convenience, powerful features, and accessibility in an integrated, affordable package. For more information or to inquire about using RIM as an individual technician or deploying RIM in your organization, please contact your Pneuma Solutions representative or visit our web site at www.PneumaSolutions.com.

Pneuma SOLUTIONS