White Paper

Honeywell

Honeywell Process Solutions

An Overview of Honeywell’s Secure Remote Access to Process Control Systems

Table of Contents

Introduction 3

Overview of Security Requirements 4

Access Control and Authentication 4

Data Encryption and Message Integrity 4

Security Measures Used by Honeywell for Remote Access to Process Control Systems 5

Virtual Private Network (VPN) 5

Client Authentication 5

SSL VPN Gateway 5

Firewalls 6

RSC Tools 6

De-Militarized Zone (DMZ) at the End-User Site 7

Honeywell Relay Node at Target Process Control System 7

Honeywell Service Node at Target Process Control System 7

Overview of a Remote Access Session 7

Conclusion 8

Appendix Architecture Drawing 9

Table of Figures

Fig. 1 – RSA SecurID key fob 4

Fig. 2 – OSI Communication Model 6

Fig.3 - Architecture Drawing 9

Introduction

There are many benefits in using subject matter experts (SMEs) to provide system monitoring, perform diagnostics and repair process control systems. The most convenient and cost-effective way to deliver these services is via a secure remote connection such that any available SME with the appropriate skills, located anywhere in the world, can provide the required service. This arrangement also allows coverage to be provided 24 hours a day, seven days a week to accommodate time zones, work schedules and available resources.

There are several reasons why the remote connection between the SME and the systems supported must be secure:

- To prevent access by inappropriate people. This includes undesirables such as hackers, but also includes well intentioned people who have no need to know or those who do not have relevant skills to offer (the vast majority of the population in the context of a process control system).
- To limit the access by SMEs to specific system nodes relevant to their skills.
- To maintain the confidentiality of a customer's plant data.
- To maintain the integrity of data as it passes from client to server and vice versa.
- To give the end-user sufficient confidence to grant remote access to a critical business asset.

By far the most convenient way to connect to a customer site is via the Internet. However, as a communications channel, the Internet is public and in normal usage it is very insecure because: (1) anyone can access it, (2) data can be tampered with, (3) data is in plain text and can be read by anyone and (4) inappropriate or dangerous software can be inserted into a system (malware). Nevertheless, the Internet is ubiquitous and very convenient; therefore the design challenge is to add features to communication channels which secure them while retaining the universal reach of the Internet.

The objective of this paper is to explain the measures that Honeywell uses to secure remote access to process control systems via the Internet and the purpose of the measures adopted. In addition, this paper focuses on security aspects rather than the functionality of the tools to provide remote services.

Overview of Security Requirements

Access Control and Authentication

It is necessary to ensure that only relevant authorized persons can access the process control system and that the scope of the access allowed is relevant to their role and skills. There are two aspects to access control:

- Knowing who the person is, referred to as authentication.
- Once the person is authenticated, limiting the access granted to that person to system components and resources relevant to their role and skills.

The concept of authentication is built around three different factors by which a person may be authenticated:

- Something they know (password)
- Something they have (access card) or
- Something they are (biometric characteristic such as a fingerprint).

At the simplest level, authentication is generally a single factor, achieved by use of a user ID and password (something the user knows). However, there are several weaknesses in this approach such as weak passwords, users not logging off, passwords being known or shared, etc. As a result, to provide more secure authentication, another authentication factor will often be employed. Such an approach is called two-factor authentication. In order to be allowed access to a system, the user must pass both challenges. Use of two-factor authentication greatly increases the probability that the person requesting access really is who they claim to be.

Honeywell uses two-factor authentication to access the Remote Service Center and then the process control system. The two factors that Honeywell uses are:

- Use of login ID and password to access a Microsoft Active Directory domain (something they know) and
- Use of a hardware authentication token, such as an RSA SecurID key fob authenticator as shown below, displays a code which changes every minute (something they have).

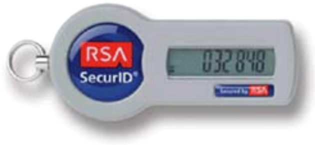

Fig. 1 – RSA SecurID key fob

Data Encryption and Message Integrity

In order to prevent users on the Internet from viewing confidential data such as user IDs, passwords or end user data, it needs to be encrypted. All forms of encryption depend upon changing the plain text data into cipher text via an encryption algorithm and one or more keys which are known only to the sender and legitimate receiver of the data. A major problem with the use of private keys is how to send the key in a secure way to the receiver. Obviously this cannot be done in a secure way via the Internet.

Secure data transmission over a public channel such as the Internet normally depends upon Public Key Infrastructure (PKI). PKI includes elements which allow the sender to be authenticated, the message to be encrypted by the sender and decrypted by the receiver, and a means of ensuring that the message has not been altered in transmission (accidentally or deliberately). PKI is a complex subject but in essence it includes:

- Use of digital certificates to authenticate the sender of the message. Digital certificates are usually issued by a Certificate Authority. The certificate confirms that the owner of a specific public key is who they say they are.
- Use of a pair of encryption keys, a public key and a private key. The sender's private key is known only to the sender and likewise the receiver's private key is known only to the receiver. Thus the problem of private key exchange does not arise. By using the keys as a pair, secrecy can be maintained. The message is encrypted using the sender's public key but is decrypted using the receiver's private key. Public keys are published via digital certificates.

In addition, measures need to be taken to ensure that the content of the message is not changed while in transit, either deliberately or accidentally (corruption).

Security Measures Used by Honeywell for Remote Access to Process Control Systems

Virtual Private Network (VPN)

The essential communications conduit used by Honeywell to remotely access control systems is a Virtual Private Network (VPN). A VPN adds a logical privacy layer over the top of another physical network (the Internet) in order to secure it from public access. When a VPN is set up, a tunnel is said to be created between computers on the underlying network. A VPN can provide authentication, encryption and message integrity.

In order to create the VPN, Honeywell uses Secure Sockets Layer (SSL). SSL includes features which provide:

- A means of encryption key exchange and hence encryption.
- A means of authentication.
- A means of protecting the integrity of data while in transit.

SSL is a complex subject and this document only provides a basic overview of its function.

Honeywell uses VPNs between the client and the Remote Service Center (RSC), within the RSC itself and between the RSC and the target control system. The remote access architecture drawing shows the various VPN tunnels used to communicate from a remote client to the Honeywell Service Node within the target control system's Process Control Network (PCN).

Client Authentication

The individual user is authenticated via a user ID and password. The user has a domain account and a hardware authentication token. These measures provide two-factor authentication and a high degree of certainty that the user really is who they claim to be.

SSL VPN Gateway

The SSL VPN Gateway acts as a single portal through which all clients access all applications within the RSC. There are two halves to the SSL VPN Gateway which intercommunicate via a special communications channel. The communication from the client to the portal uses the Internet Protocol (IP) which is then encapsulated within an SSL tunnel. This communications conduit terminates on one side of the SSL VPN.

The SSL VPN Gateway then sets up a separate communications conduit from itself to the RSC's applications server. This conduit also uses the IP which is then encapsulated within another VPN tunnel. The two communications channels then inter-communicate using a Content Intermediation Engine. See the diagram (Fig. 2) below. The SSL VPN Gateway will parse the input received from the remote client and rewrite it to the RSC application server. This technique prevents the propagation of network worms from the client to the RSC or vice versa.

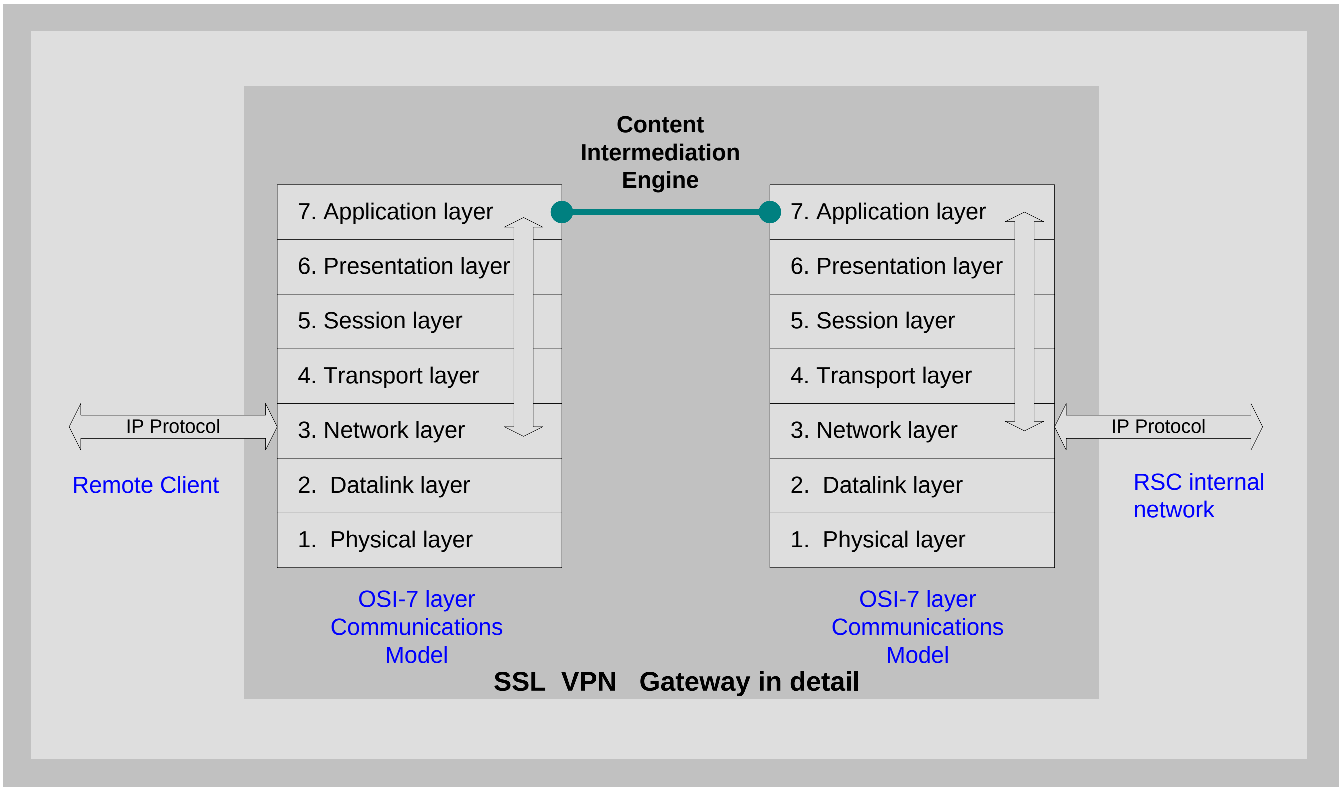

Fig. 2 – OSI Communication Model

Firewalls

In simple terms, firewalls are a means of blocking or allowing data transmission between computers on a network on a conditional basis. They provide means of filtering messages by:

- IP address
- Protocol
- Direction of data flow and
- State of the communications channel (eg who initiated the communications). This technique is called stateful filtering.

Modern firewalls may be very sophisticated and may include numerous additional features beyond the basic ability to filter messages.

Honeywell makes extensive use of firewalls within the RSC as can be seen in the architecture drawing. In addition, Honeywell strongly recommends the use of firewalls at the end-user site to protect the DMZ and the process control network.

RSC Tools

The tools used in the RSC are hosted on various servers. For the purposes of this paper, the main servers are the application server, the database server and the communications server.

Used as a set, these nodes provide the tools which are used to monitor and diagnose the target process control systems. Each of these nodes is protected by a firewall.

In addition, there are other servers such as an antivirus update server, a patch server and a test system. These servers are used to support various remote update services offered by Honeywell.

De-Militarized Zone (DMZ) at the End-User Site

A DMZ is a separate network which acts as a buffer zone between two physical networks that do not trust each other but need to exchange limited amounts of data. The DMZ is created by use of a firewall between each of the networks and the DMZ, an intermediate layer between the two untrusted networks. Honeywell always recommends that if a PCN needs to connect to another network (which is not trusted by the PCN), that this connection is made via a DMZ.

Honeywell Relay Node at Target Process Control System

The Service Node is a server provided by Honeywell which is connected to the DMZ of the target process control system.

As can be seen in the architecture drawing, the RSC sets up a VPN tunnel to the Service Node. The Relay Node, which is in between the RSC and the Service Node, sets up a separate TCP/IP session with the RSC and the Service Node.

Honeywell Service Node at Target Process Control System

The Relay Node in the control system’s DMZ connects via a VPN tunnel to the Service Node located within the control system’s security zone.

Important security features applied to the Service Node (SN) itself or to the communications channels to and from the SN include:

- The SN is configured only to connect to a specific IP address within the RSC.
- All communication with the SN is over port 443 (this is the TCP port used for secure HTTP). As a result, no modification of a corporate firewall is normally needed since port 443 will almost always be open anyway to allow SSL communications to secure web servers on the Internet.
- The RSC and the remote control system do not reveal their true IP addresses.
- The SN controls the access allowed to the control system nodes from outside the PCN security zone. The end user is always in control and decides what access is allowed per user and can terminate a connection at any time if desired. In addition, the end user can approve each remotely initiated activity and can supervise the remote operations.
- Communication from the SN is initiated outbound only. This approach allows very effective use of a stateful firewall since all legitimate communication is initiated from the SN. Thus any unsolicited communication from outside the PCN to the SN is illegal and will always be filtered. If the RSC requires communication with the SN, it sets a request flag within its communications server. The SN routinely polls the communications server. If, during a poll, it sees a pending request, it will validate that request versus its stored security policies and only grant the request if it is within policy. The policy limitations may refer to the nature of the request itself or to the user requesting the access. See section 3 for description of the setup of a remote access session to the SN.
- The Service Node and the communication server in the RSC authenticate each other using PKI before initiating communication.
- When it is required to download files, such as antivirus updates or patches, this transfer is “pulled in” toward the control system. Access control is arranged such that unsolicited “pushing” of files from outside the PCN security zone is prohibited.

Extensive session recording and logging facilities are also provided within the Service Node to provide a comprehensive audit trail of accesses and activities.

Overview of a Remote Access Session

This section will cover the steps of a typical remote access session to show how the various security measures detailed in earlier sections manifest themselves. In this scenario, imagine that an end user has requested remote support from Honeywell via telephone to investigate some problem with the control system.

1. In response to a telephone request by the end user, a Honeywell support engineer logs in to the RSC using two-factor authentication

2. The Honeywell engineer then accesses the relevant site and system. (The engineer will only be able to see systems to which he/she has access rights in the overall list of customers, sites and systems.)

3. The Honeywell engineer will navigate to the alarm and event list for that system in the RSC and examine the entries.

4. In this scenario, imagine that the alarm message in the RSC is inconclusive and the Honeywell engineer wishes to run a diagnostic test on the target control system in order to gather more information. A set of diagnostic scripts is already stored on the SN. The diagnostic test will produce some results which will be fed back to the RSC.

5. Depending on the policies agreed with the customer, remote access to the control system and initiation of a diagnostic script may be restricted as follows:

    - Remote access may only be allowed at certain times of the day and days of the week.
    - The remote access session may be limited in its duration.
    - Access may only be allowed to certain devices.
    - The end user may need to approve the remote initiation of diagnostics.
    - The end user may wish to inspect the results of remote diagnostics via a view of the remote desktop.
    - In all cases, the end user can cancel a remote access session at any time.

In this scenario, imagine that the end user is required to approve the running of diagnostics.

6. Via the RSC, the Honeywell engineer requests to run a specific diagnostic script on the target system. When the SN next polls the RSC, it will see this request. It will then check the request against its stored policies and see that the user is authorized to run this diagnostic script but only subject to approval from the end user.

7. The SN will send an email request to the end user's email server on the corporate network. The PCN firewall would be configured to allow outbound email but not inbound to the PCN. The end user engineer then receives an email which states that the Honeywell engineer is requesting permission to run a diagnostic script. If the end user engineer is satisfied with that request, he/she will login to the SN and approve the request.

8. Once the diagnostic script has run, typically results will be sent back from the SN to the RSC.

9. The remote access actions, the script running actions and the actual remote session keystrokes will be recorded for audit purposes and for subsequent replay if required.

Conclusion

Honeywell views the security of remote access to process control systems as being of paramount importance. As described here, Honeywell's security policies and processes are rigorous in order to achieve the required level of security to secure our customer's process control networks.

Appendix Architecture Drawing

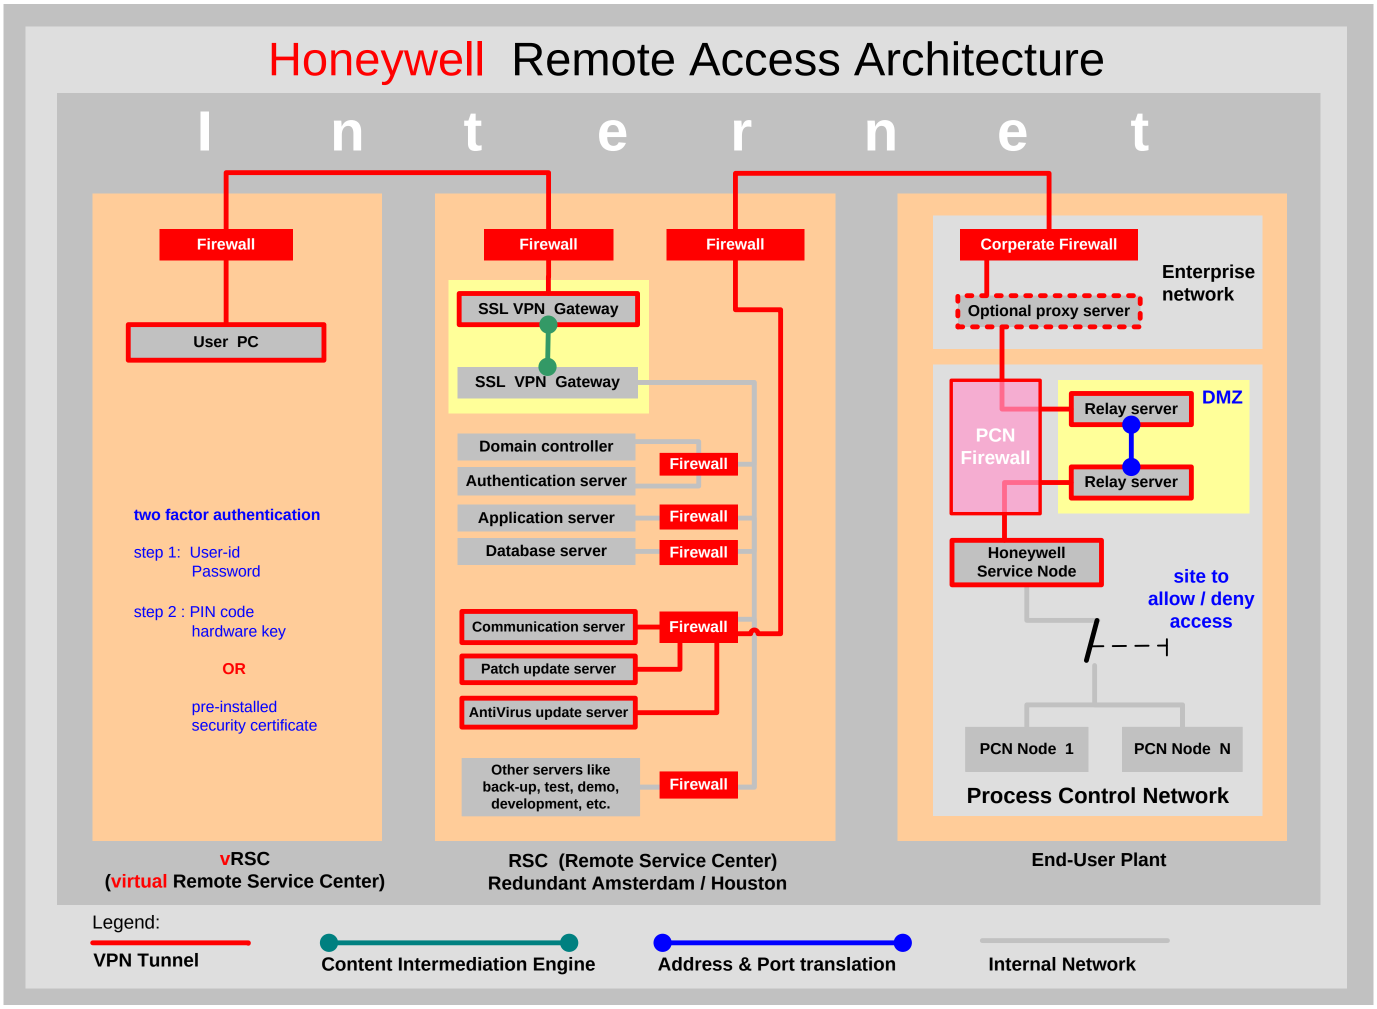

Fig.3 - Architecture Drawing

More Information

For more information about Remote Access, visit our website at www.honeywell.com/ps or contact your Honeywell account manager.

Automation & Control Solutions
Process Solutions
Honeywell
1860 W. Rose Garden Lane.
Phoenix, AZ, 85027
www.honeywell.com/ps

WP-10-09-ENG
August 2010
Printed in USA

Honeywell