

Configuring and Managing Remote Access for Industrial Control Systems

November 2010



**Homeland
Security**

CPNI

Centre for the Protection
of National Infrastructure

Control Systems Security Program National Cyber Security Division



Disclaimer

This report was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any employee, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use, or the results of such use, or any information, apparatus, product, or process disclosed in this publication, or represents that its use by such third party would not infringe privately owned rights.

Executive summary

Industrial control systems play a vital role in critical infrastructure. The requirements for their high availability and proper functioning demand that the systems be protected from both intentional and unintentional incidents that can impact their operation. In the past, the risk to these systems was mitigated by ensuring complete separation of operational domains from external networks and access to the control function was limited to authorised users with physical access to a facility. Today, business demands (such as increased and faster online access to real-time data, using less resources) has led to the rapid deployment of modern networking technologies, which has accelerated the interconnectivity of these once isolated systems. This new connectivity has empowered asset owners to maximise business operations and reduce costs associated with equipment monitoring, upgrading and servicing, whilst creating a new security paradigm for protecting control systems from cyber incident.

Part of the security equation involves how operational assets are accessed and managed and how the cyber security posture of a control system can be impacted if the management of remote access is not understood by business or is conducted poorly. However, as is the case with modern cyber security countermeasures, the application of proven and accepted remote access solutions may not map perfectly to control systems environments. Requirements for availability and integrity, combined with the unique nuances and attributes often found in 'purpose built' systems, drive new demand for guidance as it pertains to creating secure remote access solutions for industrial control systems environments.

This good practice document provides support for developing remote access solutions for industrial control systems. Common good practices from standard information technology solutions will be presented in the context of control systems environments, along with insight into how remote access solutions can be deployed in a manner to mitigate cyber risk unique to control systems architectures. The goal of this practice document is to provide guidance regarding the development of secure remote access strategies for industrial control systems environments.

In using this practice guide, no two control systems will be identical. As such, no single secure remote access solution is applicable to all possible architectures and no single remote access solution can provide adequate security without a defence-in-depth approach. However, by exercising caution and generating and implementing concise requirements based on good analysis, secure remote access solutions can be deployed and maintained.

Keywords

Industrial control systems, SCADA, remote access, role-based access control, remote connectivity, monitoring, secure vendor access, defence-in-depth, firewall, intrusion detection, encryption, demilitarised zones, security zones, policy and procedures, patch management.

Contents

Executive summary	2
Keywords.....	2
Recommended practice: configuring and managing remote access for control systems	
Introduction and definition.....	5
Remote access in industrial control system architectures.....	8
Roles and remote access in control system architectures	9
System operators and engineers for local systems	10
System operators and engineers for remote systems	10
Vendors	10
Integrators and system support specialists.....	11
Field technicians	12
Business partners	12
Customers	13
Supply chain operations.....	13
Managed service providers.....	13
Other considerations.....	14
Unintentional entry points	15
Remote access in modern ICT architectures	17
Requirements	17
Types of remote access solutions	18
Security Considerations	22
Determining requirements.....	23
Attack vectors	23
Secure communications.....	25
Access control.....	25
Logging and reporting	26
Network architecture security.....	28
Improving security.....	29
Remote access in control systems architectures	32
Remote access security considerations unique to control systems	33
Applying good practices	35
Periodically undertake formal threat and risk assessments.....	36
Eliminate all direct connections.....	37
Secure modem access	38
Create a physical and logical DMZ separating corporate and IT environments	38
Create separate authentication servers for separate roles (vendors/integrators).....	39
Enforce a security assurance policy for all remote access	40
Full tunnels.....	40
Password policy	42
Two form factor authentications.....	43
Authorisation levels.....	44

Dedicated client hardware and software	45
Session termination	45
Managing remote access in control systems environments	47
Case study: Marstad municipal operations.....	48
Background.....	49
Control system architecture	49
Risk factors	51
Solutions for securing the municipality's remote access.....	53
Conclusion	56
Glossary	
Acronyms	58
Nomenclature	59
Figures and tables	61
Further reading	62
References.....	65

Recommended practice: configuring and managing remote access for control systems

Introduction and definition

Information infrastructures across many public and private domains share several common attributes regarding information and communication technology (ICT). This is particularly true in the industrial control systems domain, where an increasing number of organisations are using modern networking to enhance productivity and reduce costs by increasing the integration of external, business and control system networks. However, these integration strategies often lead to vulnerabilities that can greatly reduce the cyber security posture of an organisation and can expose mission-critical industrial control systems to cyber threats. The opportunities for enhancing business operations are seemingly endless, with one of the major advantages being the ability to increase the command and control function by leveraging remote access capabilities. Without applying appropriate security safeguards, remote access functionality can create opportunities for cyber adversaries wishing to cause harm or damage to critical processes that may seriously affect the lives of people, health, society, the economy and the environment.

This document provides guidance for developing secure remote access strategies for organisations that use industrial control systems. This document is to be used in developing or updating strategies related to managing remote connectivity between operational assets, peers, vendors, operators and other elements that require access to critical information, devices or process data.

From a definition perspective, this document will assume that remote access is defined, in its simplest form, as ‘the ability for an organisation’s users to access its non-public computing resources from external locations other than the organisation’s facilities.’¹ To extend this to control systems, consider that remote access also includes ‘accessing data, a system, or a service inside a physically and/or logically protected network from a system or device outside that network.’ Combining the two, the definition of remote access for this practice guide is:

‘The capability for an organisation’s users and operators to access its non-public computing resources, data and systems that reside inside a physically and/or logically protected network from external locations that may be considered outside that organisation’s network.’

This definition is useful in many circumstances,^a in particular in control system domains and has several security elements that are applicable to this practice guide:

- It provides allowances for the fact that a single operator may have administrative oversight over several disparate systems that are considered to be within an organisation’s information enclave.
- It implicitly acknowledges that remote access can be interrupted, prevented, captured or hijacked through deliberate actions of a separate party without that party having to circumvent physical or logical security controls, even when the communications are travelling across media owned and operated by the data/system owner.

a. A more rigorous definition is ‘data communications from one network enclave to another network enclave; where a network enclave refers to a networked group of one or more systems partially isolated from other systems using some protection mechanism, either logical, physical, or both.’

- It can exclude communications within physically protected areas, such as compounds or buildings, reducing the scope of the discussion (security of communications within physically protected boundaries is still an important issue, but should be considered separate from and beyond the scope of remote access).
- It includes all communications over equipment whose physical and logical security cannot be validated explicitly by the organisation using the equipment in an area of communications security that traditionally does not receive enough attention, even today.

This definition expands the scope of remote access to include examples that traditionally have not been considered 'remote access.' For example, connections between geographically disparate sites using private third-party telecommunications lines are not normally considered remote access, despite the fact that the telecommunications lines are owned by a separate telecommunications company and leased by the organisation looking for a communication mechanism between two sites. However, they are included in this definition. The definition encompasses long-range communication channels, such as fibre or microwave, where the equipment is owned by the data or service owner, but the physical security of the transmission media is outside the direct control of the data owner.

Remote access security functionality and features help create electronic pathways to grant authorised and authenticated access into a trusted network from a location that would otherwise be considered untrusted. In this definition, the non-public (trusted) network would be considered the control system network.

Although this document is titled *Configuring and Managing Remote Access for Control Systems*, the material is intended to be applicable to any architecture involving industrial control systems, process control systems, Supervisory Control and Data Acquisition (SCADA), or distributed control systems. The term industrial control systems is to be considered a general term applying to all these system types sharing similar characteristics and is in line with the definitions used by the contemporary communities of interest and other standards bodies.


Background

The critical infrastructure systems that support major industries, such as manufacturing, water, transportation and energy, are highly dependent on information systems for their command and control. While a high dependence on legacy industrial control systems still exists, critical industrial control systems are migrating to new communication technologies. As a result, common communications protocols and open architecture standards are replacing the diverse and disparate proprietary mechanics of industrial control systems. This replacement can have both positive and negative impacts, especially in the areas of system integration and support. While the traditional isolation of the control system demanded onsite maintenance from integrators and vendors, modern communication mechanisms can facilitate remote connectivity from almost anywhere. In addition, this new interoperability provides for operators and asset owners to allow their own disparate resources to remotely connect to their control systems from anywhere on an as-needed basis. As is the case with any contemporary ICT architecture, the cyber risk is proportional to the security countermeasures deployed to protect against unauthorised remote access.

The protocols and communication standards that provide increased interoperability in the industrial control systems community build on and use, in many cases, the same technologies that have been exploited and compromised on the internet and corporate networking domains. The same is true for those technologies associated with remote access and this can often

create situations that cause industrial control systems to inherit undesirable security vulnerabilities. Research indicates that the mitigation strategies used in contemporary ICT systems may not always align perfectly to the industrial control systems domain. The unique nuances associated with both availability and integrity within control system architectures require that contemporary security countermeasures be deployed with a specific purpose.

Figure 1 provides a notional diagram that illustrates the traditional isolation of a control system environment from supporting corporate architectures and peer sites. This simplified view showcases that access to the control system domain, be it for system operation or system support, was either by physical access to the facility or remote access via telephonic means (modem). This illustration is a very good approximation of traditional legacy architectures and showcases traditional vendor access provided by a dial-in capability. Onsite support, which was often the norm, would tend to be expensive due to time and materials and organisations looked for ways to reduce costs wherever possible. With regard to security, access policies and controls were developed based on the assumption that external environments contained threats that were both malicious and hostile.




(Figure 1: Traditional isolation of corporate and control domains)

But as interoperability also provided solutions for cost reduction and ease-of-use, access into the control systems environment was extended beyond simple modem and physical access. As organisations grew, the requirements related to reporting and business oversight also expanded and continued to expand. This drives the need to facilitate operational information to different areas of the business and that could include regulatory authorities, remote operations, peer

sites, partners and even application and hardware vendors as well as third parties providing services.

Remote access in industrial control system architectures

Figure 2 shows an integrated architecture that has connections from external sources such as the corporate local area network (LAN), peer sites, vendor sites and the internet. This graphic also illustrates the concept of an external communications infrastructure, an element common in control systems architectures, be they localised or disparate across large geographic areas. It is this external communications infrastructure that supports connectivity to elements of remote operations, remote facilities, business partners and vendors. The external communications infrastructure could also be considered a connection point to the internet, the current de facto mechanism for remote users to gain access to business or control system operations (i.e. telecommuting).



(Figure 2: Integrated networks)

From figure 2, the integrated architectures, if compromised, clearly could provide an attacker with various avenues for accessing critical systems, especially if remote access mechanisms are taken into consideration. The issue of convergence raises concerns regarding how attackers could create vectors into trusted control architectures simply by compromising trusted resources in remote operations, remote facilities, remote business partners and even vendors. These

issues clearly showcase the importance of creating effective and secure remote access solutions when dealing with mission critical control systems. But as discussed above, the implementation of secure remote access strategies can in many cases be nontrivial as the requirements regarding availability and integrity make standard solution deployments impractical and ineffective. More importantly, the deployment of a secure remote access solution may not only increase the cyber risk profile of a control system environment but may negatively impact the necessary attributes of high availability and data integrity.

Roles and remote access in control system architectures

Because securing remote access is an integral part of any defence-in-depth strategy, the foundation of creating usable guidance as it pertains to control systems environments must include both users and the technology to be accessed remotely. To generalise control system architectures is difficult and to develop a recommended practice for securing remote access that is applicable to all architectures is impossible. The uniqueness and diversity of both vendor and purpose-built systems create a landscape of diversity that simply cannot be addressed with a single solution. However, common elements, such as users, roles, existing technology and architecture types, can be reviewed and their attributes can be leveraged. It may help organisations to shape their remote access strategy by determining who requires access to certain resources as well as understanding attack vectors that can be created unintentionally.

Understanding both users and roles can have a significant impact on how the remote access strategy evolves. In most control systems operations, the roles that would require remote access to control assets may include, but are not limited to:

- System operators and engineers for local systems
- System operators and engineers for remote systems
- Vendors
- System integrators
- System support specialists and maintenance engineers
- Field technicians
- Business partners
- Reporting or regulatory entities
- Customers
- Supply chain representatives
- Managed service providers

Developing a list that is complete for all systems in all sectors is not possible and the list should be augmented as needed. The roles of the users that would require remote access to mission-critical operations can be extensive and the assignment of specific access depending on those roles can be complicated at best. The assignment of remote access roles and credentials is expected to be embedded in the organisational cyber security policy supported by the remote access methodologies discussed in this paper. Here, some of the more common roles are reviewed in developing a control system remote access strategy.

System operators and engineers for local systems

The people who have the most need for access into a control system environment are going to be the regular system operators and the system maintenance engineers. The requirements for remote access to local systems are best understood when the growth or expansion of a control system will require an operator or engineer to have oversight over disparate (but connected) resources. Modern networking does not always facilitate connectivity that automatically creates access control permission structures and if it does, it certainly will incur a cost. Asset owners rely on the concept of flat networking to allow operators and engineers to have both local and remote access to mission-critical systems.^b Local systems, by definition, include information resources that are not actually situated on the control system network. But, they are located on conjoined domains that are, for example, responsible for aggregating operational data or preparing content for a customer web-based service.

Access to and from critical control system assets in the modern environment is usually LAN-based, but still should be considered remote if the operator is traversing across different networks. Virtual Private Networking (VPN) is often considered the best approach in securing trans-network communication. Surprisingly, many organisations feel that these countermeasures are often unnecessary due to the trust relationship that exists between operator consoles and services requiring access. Trust relationships inherent in most control system domains not only involve operators but also field device technologies. The foundation of control system operations incorporated exclusive trust between operators and field devices, which was a capability well suited for completely isolated networks. Unfortunately, this trust has been carried over into modern control system architectures and this inherent trust is easily exploited by an adversary that has managed to compromise the control system or administrative networks.

System operators and engineers for remote systems

As control systems environments continue to grow both in terms of size and capability, asset owners will be faced with the challenges of maximising operations while managing expenses related to personnel. These requirements will introduce demands for an organisation to have an operator administer more than one system in order to optimise operations. More importantly, asset owners need to consider the emerging complexity of distributed environments that facilitate single operators having responsibility over many disparate elements. In many cases, the operator will not need the same levels of authoritative access across all environments. Obviously, this can impact the access profiles an engineer can have when managing operations across several system domains.

Vendors

Traditionally, asset owners were not concerned about remote access activities from vendors. When systems were completely and totally isolated from any external connectivity and the only access to the network was either physical or by modem, entities felt comfortable with the risk associated with vendor remote access. In modern control systems environments, vendors often deploy their technology complete with remote access (e.g., a modem) embedded into the

b. To enhance productivity, it is not uncommon for asset owners to assign processes and procedures specific access capabilities across flat networks.

solution to expedite timely support operations, operations that can include system restoration, system upgrades and performance monitoring. Vendors can perform these functions as part of a larger support services contract. In many cases, these contracts demand that remote vendor access be available and not necessarily include operator interaction.

The impact of vendor access on control system security posture is interesting. The electronic security perimeter associated with protecting ingress and egress data streams becomes very ethereal and hard to qualify. Whereas the traditionally isolated model allowed the asset owner to clearly define the security perimeter, the protection of the information enclave is now the responsibility of the vendor with the remote access. Technically, all the points of presence the vendor has on the internet could conceivably expose the asset owner, thereby changing the risk profile of the control system dramatically as soon as remote access is obtained. This issue is the cornerstone of the argument that control system solution vendors must provide adequate security countermeasures for remote access into the systems they are supporting and that the historical trust relationship shared between a vendor and an operator can no longer be assumed to be safe and uncompromised.

Asset owners that rely on remote access from vendors to support their operations should consider a number of different factors in developing their overall remote access strategy. These can include, but are not limited to:

- The procurement guidance language (see Reference 1) may be used to agree on the establishment of a guaranteed secure remote connection between the vendor and the owner's control system.
- Asset owners should be aware that the vendors could be exposed to untrusted and hostile environments. A compromised vendor network could enable an attacker to piggyback on a trusted connection into the control system.
- Vendors will usually require remote access for two reasons: emergency operational support and system maintenance. Asset owners should recognise that the latter of these two can be scheduled and definitive protocols for remote access connections can be established and monitored.
- Asset owners need to be aware of the possibility that a vendor can deploy their remote connectivity solutions (i.e., modems) in a standardised manner and as such the authorisation credentials used to access a large number of control systems may be very similar if not identical.

Integrators and system support specialists

Generally speaking, integrators differ from vendors in several ways. With regard to remote access, the most critical difference is the presence of an integrator in solution provisioning that often prevents the asset owner from directly interacting with the vendor. Although it is not uncommon for an integrator to also be a vendor, this single degree of separation can greatly influence the cyber risk posture of an organisation. Indeed, the remote access issues relevant to vendors may also be applicable to integrators.

Security carries a nontrivial cost associated with it. Some integrators are unwilling to take the necessary steps to ensure secure remote access above and beyond activating default countermeasures embedded in the remote access technology. More often than not, integrators will have a standardised solution for remote access that usually involves either a modem or some other economical means of providing support from an offsite facility.

Field technicians

One of the emerging areas is how field technician's access critical systems either for operations, maintenance, or the provisioning of performance metrics. As many sectors begin to deploy programs that ensure rapid collection of field operations data, the methods used for remote connectivity to facilitate this data transfer requires investigation. Traditionally, technicians that operated in the field usually provided operational data at the end of the shift, usually by physical means and at a central location. New requirements to support advanced business needs, such as those that relate to grid operations or drinking water safety, demand that technicians are in constant communication with engineers administering control system operations.

Many asset owners use modern communications technologies to facilitate field technician's requirements to submit data to utility management systems. In many cases, these management systems are contained in the corporate or business environments even though the data are used to optimise control systems operations. Some methods for the remote communication, which allows for the timely data exchange between business and field technicians, need to be reviewed from a security perspective, because the compromise of these communication channels can provide an adversary with significant access to core system operations. In addition, as field data are a critical component in meeting the demands for high availability, remote access for field technicians must also be maintained to allow access as needed.

Asset owners supporting sectors that require constant support from field technicians have several opportunities to leverage modern secure remote access solutions. In addition to providing field technicians access to remote facilities that share a trusted network, remote access solutions involving virtual networking and multifactor mutual authentication to centralised servers can provide a significant amount of protection of critical communications paths while maintaining business objectives.

Remote access policies for field workers also must include consideration for physical security. Organisations need to assume that the computing resources used by field technicians will be targets and an increased risk exists pertaining to field technician computing devices being stolen. If the remote access solution for the organisation empowers the user of the field computer to gain access to control devices, then it should be assumed that a thief or attacker who is in physical possession of the stolen computing device may also get access to control system assets unless proper security measures are taken.

Business partners

The risk associated with business partners has always received attention and whether the relationship between an asset owner and a business partner is young or old, there is a constant concern about information theft and the exfiltration of proprietary data.

With the advent of modern network connectivity, the accessibility available to business partners to extract information from peers has grown considerably. This accessibility, unfortunately, is often granted by default based on the historical relationships between business partners and asset owners.

With a business partner having access directly into a control system environment, asset owners must take responsibility for ensuring that this access is managed and secured. In addition to providing dedicated lines and cyber security policies associated with the accessibility a business partner can have, constraints limiting the business partner's access to control system operations need to be deployed.

Although it becomes very easy to understand the consequences associated with a cyber attack from a hacker or a disgruntled insider, sometimes organisations fail to appreciate the consequences associated with a critical control systems failure that is a result of poorly managed remote access from a business partner. In addition to impacts that can, for example, range from failure in a batch mixing process to an event causing death, cyber incidents involving business partners can have catastrophic impact that could (in a worst case) lead to the extinction of the business itself. When considering remote access solutions, situations that demand business partners to have access to mission-critical operations require a fresh look from a consequence analysis perspective. The output of that analysis can provide insight to appropriate controls associated with a remote access capability.

Customers

Although the provision of access to customers into control systems environments has historically been quite rare, recent developments in data aggregation methods and rapidly emerging service requests from the customer community are changing this. Using ideals from smart grid innovations, with the end state aiming to empower users with choices regarding better energy usage (and in turn allow utilities to make better choices for energy management); the traditional separation between end users and critical control system operations is becoming blurred. In essence, the separation between critical energy distribution operations and the customer home may simply be a utility network and a farm of web presentment servers. Taking this simplified approach, it is clear that the remote access provided by asset owners to their customers requires a substantial review from a cyber security perspective.

Supply chain operations

Supply chain operations is often one of the most overlooked elements when considering remote access solutions. In many modern business environments where control systems are involved in manufacturing, the supply chain is responsible for ensuring the timely and accurate delivery of raw materials and services. Many asset owners in the manufacturing domain, as well as some in the utility space, would be unable to succeed if not for the provisioning of materials and support from their supply chain. Because the success of a business operation is tied to the speed at which they can create products or service, it is critical that supply chain operations have access to information that allows for the timely replenishment of materials. Often this access involves direct communication to control systems environments, with particular focus on monitoring stock levels and production outputs. Supply chain operations do not just impact the business from a pre-manufacturing phase but also provide post-manufacturing support with things like quality assurance, transportation and safety monitoring.

Managed service providers

Asset owners are including the capability for managed service providers (traditionally only used within ICT domains) to have access directly into control system operations. Interestingly enough, these services are often presented from the perspective of cyber security and the service providers require access into the control system to assess security postures. These services are also sold by the vendor community, where the vendors are given full-time access to the control systems environment to ensure operational stability.

From a financial perspective, it makes good business sense to outsource system operations to a trusted third party, offsetting the cost associated with having to staff certain operational

requirements. If the managed service provider can augment traditional support services with services that support integrity and availability of the system, the benefit could greatly outweigh the risk. That said, the remote access provided to the managed service entity needs to be carefully crafted under an extended security agreement and the burden of responsibility for securing the communication path lies with the service provider. Although the entity could help in establishing the remote-access function of the management and protection of that remote access pathway, security needs to be cared for by the provider.

The emerging interest in managed security service providers having access to control systems environments raises concerns in many areas, especially those that involve the service provider actively trying to mitigate threat in production environments. Often, many service providers are granted full authority to actively try and mitigate cyber incidents using their remote access capability. In many circumstances, this could be a plausible and desirable activity, but the sensitivity of and unique nuances associated with control systems environments creates a situation where the access can have adverse impact. A simple example of this would be the introduction of antivirus upgrades into a production environment or the mitigation of vulnerabilities associated with control system peer-to-peer communications. In addition to ensuring the access path is limited to only the managed service provider, there can be significant negative impact on production environments because of the implementation of untested security patches and mitigation strategies.

Other considerations

Asset owners required to provide remote access to any of the elements discussed in the previous section will have a number of options available, including tunnelling; providing direct access to applications, access portals; and remote desktop access. Considering the availability and integrity demands usually required of control systems, asset owners must be rigorous in ensuring that the remote access solutions are balanced appropriately with business requirements. When considering the options available for provisioning remote access, organisations should be cautious so that no unintentional entry points are created when during implementation. Regardless of the solution, several common elements are pervasive across all remote access technologies:

- Remote access allows users to store critical information locally on their computer or device.
- Remote access solutions are not restricted to using single modes of authentication. The risk associated with information disclosure or compromise can sometimes demand several modes of authentication combined with several different modes of server access.
- Cryptography has and will continue to be part of the remote access solution, but cryptographic communications may impact the timeliness of communications expected and the processing capacity of control system elements within some critical operational environments.
- All remote access solutions depend on the physical security of the devices and authentication elements (e.g. passwords, tokens) initiating the remote connection.

In establishing the different types of roles required for access into control systems domains, the risks are unique to industrial automation environments. Organisations cannot be expected to

develop a secure remote access strategy without fully understanding some of the complexities associated with high availability networks that are often composed of both modern and legacy-type architecture elements. Issues, such as the interconnected IP and serial networks, as well as other communication mechanisms that are naturally devoid of security countermeasures, play a significant role in prioritising issues. Generally, several issues should be considered as part of the foundation used to create control system remote access solutions.

Unintentional entry points

The convergence of control systems with other architectures has historically occurred at a rate that allows for unintentional and unsecured access to survive. Technological capabilities that were commonplace in isolated systems often remain in updated environments either because the removal of the technology apparently impedes productivity or the technology needs to be maintained to meet cultural requirements. Two of the more common examples of this are data servers and printers, both of which tend to be network enabled and are often considered benign from a cyber risk perspective.

Most common environments today have updated modern server and printer solutions and these complex machines possess considerable computational horsepower that can easily be leveraged to support various activities, including those that benefit a cyber attacker. Printers have and continue to be considered a vital element in control system operations and are used for network diagrams, productivity reports and other operational functions that require printed hard-copy materials. Printing devices are usually accessible by everyone by default and often are deployed as a trusted information resource that can connect to more than one network. This is common practice, as the risk element of the printer has usually been considered low. Significant cost savings occurs in having only one printer serve several networks. But, with the modern printer having all the capability of a modern desktop computer, the compromise of it (compounded by the access it has across conjoined domains) can provide an attacker with significant access into the operational environment.

Unintentional entry points can also be created with the deployment of cost-saving interoperability solutions such as wireless. The entry points into the operational domain need not necessarily come from the control system directly. Rather, the compromise of a wireless connection that provides trusted access into an automation environment can have an equally negative impact on an organisation's cyber risk stature. The countermeasures associated with securing wireless networks often extend beyond the default standard recommendations as provided by the vendor. Because several different types of wireless solutions exist, such as those for business, control operations and data sharing, organisations tend to limit the amount of rigor that is required to protect wireless communications environments where availability is of utmost importance.

The advent of modern networking technology has provided control systems operators with new capabilities that increase their effectiveness in managing day-to-day operations. Vendors have been able to meet the needs of asset owners by incorporating web-based capabilities in control systems and field devices, allowing operators to manage many assets over a broad area. However, a rapidly emerging problem exists in how unused or mismanaged communications functionality can be leveraged by an attacker, especially when the asset owner either doesn't know the functionality exists or the default deployment demands that the capabilities are enabled. Embedded systems in field devices are becoming exceptionally powerful and new capabilities, such as remote monitoring, system diagnostics and firmware updates, can now easily be executed from a distance. This can create unintentional entry points into a system that, if left unguarded, can create an access vector for an attacker.

The creation of unintentional access pathways by a remote access solution is a critical concern. Organisations need to be prepared to deploy intrusion detection and intrusion prevention mechanisms to mitigate the risk. Fortunately, a long history of successful ICT centric remote access solutions can be leveraged for control system operations. The success of the deployment becomes a function of how those traditional solutions are tuned for the control systems environment and how the cyber security mitigation strategies are appropriately balanced with business requirements.

Remote access in modern ICT architectures

In this section, traditional remote access for business networks is described, as are the techniques that have been developed to secure corporate remote access.

Requirements

To properly understand the foundational requirements for establishing remote access programs, assessing the requirements, resources and critical information and how information is accessed is critical. In the control systems environment, access into and out of mission-critical environments may not always be limited to automation processes. In many cases, legitimate elements of the entire business operation infrastructure require access to data in resources involved in control systems operations. To help shape what these requirements look like, remote access requirements are viewed from an ICT perspective.

The following list presents some of the more common priorities that traditional ICT infrastructures concern themselves with regarding remote access. The solution must provide easy and meaningful^c access to critical information and essential services to those who require it to perform their duties.

- The solution must ensure that resources^d can access their information and essential services from any location if required and allowed by policy.
- The solution must ensure that resources can access their information and essential services from multiple device categories.
- The solution must ensure that critical information is available for use at a moment's notice (according to expected service level), that the critical information is always accurate and complete and that confidential information is provided only to those who require it.^e

Critical, confidential information requires appropriate protection from abuse or disclosure. When examining data sets, security personnel consider the confidentiality, integrity and availability requirements of such data. These security requirements are always dependent on the business impact of failing to maintain them for the data set under examination. As has been discussed, cyber risk is a function of consequence and this may have a different meaning among sectors using industrial control systems.

c. Meaningful access in this sense refers to the information that is generally limited in nature and is presented to the user's own and need-to know basis. This principle of least privilege is a recommended security practice and is usually found when developing strategies to disperse information to users accessing systems remotely. Although a recommended best practice, the security countermeasure is (unfortunately) not always deployed in production environments.

d. The user is reminded that a resource can be a person or a service and that remote access solutions need to include consideration for both people and computing assets that need to share information across domains.

e. This requirement has specific applicability to control systems environments. However, as it pertains to mission-critical control systems (especially those supporting critical infrastructure), in many cases the elements of availability and integrity supersede the requirements for confidentiality.

Types of remote access solutions

As mentioned earlier, all remote access solutions are essentially the same at their core, as they provide a connection over a distance between a user (or system) and a system (or information asset). Just as there are differences in access requirements, there are technological differences between solutions. Some organisations use third-party telecommunication links and rely on the service provider's ability to secure communications over their infrastructure without question. Examples of this include leased lines, Multi Protocol Label Switching (MPLS), Frame Relay, Satellite and Plain Old Telephone Service (POTS) modems. Other entities use infrastructure owned by the communicating party such as microwave links, long-range fibre optics and copper lines. Interestingly, the technologies are the same in this example as in the previous example; the only difference is they are owned and operated by the communicating party. Still others use existing computer networks and tunnel connection between locations such as VPNs, telnet, secure shell (SSH) and others.

Modern application technology has blurred the definition of remote access in recent years. As both people and services can require remote access, the suitability of technological solutions becomes hard to discern. Many communications protocols are being encapsulated in web traffic and web-based applications are managing increasingly complex communications inside HTTP and HTTPS channels.[†] Further, internet-facing web-based applications are becoming universally accessible; yet, they manage the security of the data that they are presenting to various user communities, in some respects the data are being brought to the consumer.

Table 1 shows a sample of different remote access solutions that have been used over the years in the ICT domain that has found applicability in control systems environments.

f. These communication channels are often embedded by default in modern field device equipment.

(Table 1: Different remote access solutions)

Basic Method	User	Target	Technology	Description
POTS or Integrated Services Digital Network (ISDN) Dial-Up	Single User System	Direct to Target System	Modem	A single user has a modem connected to either a POTS, or ISDN line and directly calls a computer system, which also is connected to a POTS or ISDN modem. This is a direct connection, which usually only allows a single user at a time to gain access to a single system. This method of remote access was very common prior to widespread use of the internet and is still in use today for accessing legacy systems.
POTS or ISDN Dial-Up	Single User System	Remote Access Server (RAS) and network of systems it serves	Modem	A single user has a modem connected to either a POTS or ISDN line and calls a RAS. This RAS does not provide application services to the user. Rather, it is a gateway to other systems that do provide the user with services and authorisation can be done at any number of access points. This method extends network connectivity to the user. The RAS server becomes, essentially, a router on the network, which has a network on one side and a single endpoint on the other.
POTS or ISDN Dial-Up	Site LAN and associated network of systems	RAS Server and network of systems it serves	Modem	A packet routing device uses a modem to connect to either a POTS or ISDN line and calls a RAS. This RAS does not provide application services to the calling device. Rather, it acts as a packet routing device (or bridge) to allow network traffic to pass from one network to the other. The phone line becomes a wide area network (WAN) link and packets are routed across it as needed. This method extends network connectivity between two arbitrary-sized networks, just like any other WAN network technology. In this example, the distinguishing factor is that the connection is usually temporary - established on demand and torn down when not in use. This is not a requirement and permanent WAN connections have been established over POTS and ISDN. However, this was the exception.

Basic Method	User	Target	Technology	Description
WAN Links	Site LAN and associated network of systems Single User System (rare)	Router/bridge and network of systems behind it	Varies	A WAN link is any telecommunications method to encapsulate network traffic and transport it over long distances. In the previous example, this was accomplished using POTS or ISDN lines from a telephone company. Other long-haul communications links are available. Some are wired such as leased Telco circuits (T1, T3, DS0-3, Frame Relay, ATM, MPLS, Broadband cable, DSL). Others are wireless (802.11, WiMax, Satellite uplink, other line-of-sight microwave, short-wave). Single user and point-to-point links using these technologies are possible, but very rare because of expense. In this example, the distinguishing factor is that these links are mostly permanent, rather than established on demand.
Local Wireless	Single User System	Wireless router and the network of systems behind it	802.11	A single individual with a wireless access card, or wireless local area network (WLAN) card, is connecting the one system to an organisation's wired network through a wireless access point. The wireless access point acts as a bridge, allowing the wireless-enabled device to connect to the local LAN, as though it were connected directly using a network cable. This connection method is not yet widely considered remote access. However, the wireless signals are not bound to remain inside a space under the physical control of the network owner. Therefore, it has all the vulnerabilities of any other type of remote access and should be included in remote access technology.
VPN (Virtual Private Network)	Single User System	Remote VPN Concentrator and the network of systems behind it	IPSec, SSL-VPN, L2TP, SSH Tunnelling	A single user applies VPN technology to create a layer of protection for the user's communications and then sends that communication over any of the previously mentioned links. The VPN technology essentially creates a tunnel, where the data inside the tunnel are protected using cryptographic techniques. It can be considered a virtual WAN link, which is encapsulated within typical WAN traffic. A VPN link is always dynamic in that it is established on demand and torn down when no longer desired. VPN links need to change their cryptographic keys regularly, so tunnels are re-established with new keys often, usually more quickly than any single

Basic Method	User	Target	Technology	Description
				session of use. However, the session time for single user VPNs is usually limited. The user will use the session for a short time, then disconnect when they no longer require the link.
VPN	Site LAN and associated network of systems	Remote VPN Concentrator and the network of systems behind it	IPSec, SSL-VPN, L2TP	A network routing device creates the VPN tunnel to another network routing device. This creates a virtual route, or tunnel, between the two devices. The routing devices use this tunnel the same way they would use any direct WAN link between the two devices. Network traffic at either location can now route to systems on the network at the other location as though they were directly connected. As with single user VPNs, these virtual WAN links are dynamic and are torn down and re-established regularly while in use. This is to keep the cryptographic keys refreshed. The session for site-to-site VPNs is also limited by their use, as with single user VPNs. Generally, they are torn down when not in use. However, they are automatically re-established any time a new communication begins that needs to traverse the tunnel.
Network Connection	Single User	Critical System	Telnet, web based administration software, ftp, SSH, Citrix, Terminal Server	This type of 'remote' connection is included where the communications are travelling across media with a lower security classification than the highest classified participant. For example, a user connecting to a critical system from the corporate LAN to manage that system is considered in scope here. This is because, just like when communicating across WiMax microwave links, systems or people can access or interrupt the communication between the participants. In this example, that includes every user and system connected to the corporate network. However, it would not include corporate users accessing their e-mail from the e-mail server. It is assumed that incidental contact by people on the corporate network would be of minimal impact to the business; thus it is not considered here.

Security considerations

Cyber security can be associated with many goals, but the goals of protecting an asset's data integrity and availability are the de-facto key reasons that organisations try and protect their control network assets. Indeed, confidentiality also requires adequate attention, but the goal of this practice is to provide remote access programmatic support for the control system community. Thus, integrity and availability remain the highest priorities. In order to accomplish this goal, two basic principles are universally applicable to aid in the preservation of data security.

The first principle is the principle of least privilege, which states, 'One must provide access to information or services only to those who need it.' Alternately, this can be inverted and stated as 'One must restrict access to data or a service to only those people or systems which need to use it'. This principle assumes or implies several things:

- One or more robust protection mechanisms must be in place to restrict access to data and services.
- The requesting users and systems can identify themselves in some fashion to whatever protection mechanisms are in place.
- The protection mechanism must be able to authenticate the identity provided to it. This is a proof mechanism, where the claimant must prove beyond a reasonable doubt that they are the owner of the identity presented for consideration.
- The protection mechanism(s) must control access to only the information sets and services that the authenticated user or system is permitted to access or use.

The principle of least privilege also requires an organisation to evaluate the access mechanisms to critical field equipment and in some cases requires a mandatory review of default remote access mechanisms. As stated before, default configurations can inadvertently create access paths to critical equipment. Organisations are always encouraged to do a preliminary risk analysis associated with embedded remote access technologies in field devices.

The second principle, which is in support of the first principle, is defence in depth. Every access vector to data or services should be considered when building defences, especially those that segregate business operations from mission critical control systems. Three essential components are available for defence in depth, each of which is broken down into a series of granular controls. They are:

- Prevention - a series of controls to ensure the security of information, systems, or services;
- Detection and response - controls to detect when prevention fails and respond accordingly;
- Security management system - a set of practices designed to continually evaluate and adjust security controls to more closely meet the requirements of a particular environment, industry, or organisation.

Determining requirements

Organisations must always understand the criticality of the system or service to which a remote user wishes to connect. A security classification system is required so that every system or information asset has a level of protection that is commensurate with the value of the asset. This ensures that the level of effort required to establish and manage the controls for a particular asset is aligned with the business impact resulting from compromise or disruption. This approach has specific applicability in the control system domain where compromise of critical assets is understood and has its roots in hazard operation management.

An understanding of resources (different types of people) the organisation is willing to assign access to for a particular system or information asset also needs to be considered.

Understanding critical asset classification, combined with who needs access to those assets and under what conditions, provides the foundation for using zone and conduit models.²


The third consideration to understand is the controls necessary to comply with the security principles discussed above. To understand the controls required, one must understand the threats and possible attack vectors. Figure 3 provides a notional diagram that illustrates attack vector access points. Now that remote access is determined to be the combination of three different topics, namely access control, secure communications and reliant communications, the components will be broken down and addressed separately, then combined to discuss second order issues.

Attack vectors

Beginning at the remote user and following the connection to the data or service, remote access can be compromised at any of the following points:

- The user or system can be impersonated to fool the target system;
- The attacker can use captured or guessed credentials to impersonate the user;
- The attacker can intimidate or coerce the user to provide valid credentials, or to perform activities at the attacker's demand;
- The user's access device (laptop, PDA, etc.) can be attacked and compromised and used to access the control system network;
- The target system can be impersonated by an attacker to fool the user and thus gain credentials or other information from the user system;
- The communication can be listened to by third parties anywhere along the communication chain;
- The communication can be interrupted or jammed;
- Communications can have data injected into them by an attacker;
- The communication can be hijacked after it has been initiated (does not rely on impersonation) or intercepted during initiation (impersonate both user and target, also known as a man-in-the-middle attack);
- Parts of a communication can be replayed to a target, even if the attacker cannot decipher the content (also known as a replay attack);

- The target communication software listening for requests can be attacked and potentially compromised;
- An attacker can impersonate a valid communications node and gain access to the underlying communications medium;
- Denial-of-service attack to authentication server (e.g. radius server or RAS);
- Denial-of-service attack to outward communication device (e.g., modem bank, outside router for remote access).



(Figure 3: Attack vectors when using remote access)

Secure communications

Communications media are not generally secure. Anyone with access to electronic equipment can eavesdrop on unprotected communications anywhere along its physical route. Wired communications are susceptible to taps on the wire and wireless communications can be captured with a simple antenna, protocol analysers and other technologies. More sophisticated attackers can mimic part or all the communication, inject false data in the communication stream and even assume the identity of either or both of the communicating parties during the middle of a communication. Though technically more difficult, even fibre optics are vulnerable to taps, interruption and forging.

These vulnerabilities require several preventative controls to remedy them.

- **Secure Channel** - a mechanism to prevent a third party from capturing the communication in an intelligible format. This is accomplished using advanced cryptographic capabilities.
- **Robust Channel** - a mechanism to ensure the integrity of the communications so that the communication cannot be altered in transit without detection. This is accomplished using advanced cryptographic capabilities.
- **Available Channel** - a mechanism to ensure the availability and timeliness of the communications so that a reliable data channel can be established. This is accomplished using various anti-distributed-denial-of-service techniques, load limiters and anomaly detection.
- **Device Hardening** - mechanisms to enhance and ensure the security of the systems at either end of the communication, so that they cannot be compromised. This includes robust access controls (see next section), patch management, configured according to the principle of least privilege, etc.

Access control

Systems that must protect some information and services, or must allow only certain people or systems access to information and services, require access control capabilities configured in compliance with the principle of least privilege. Traditionally, ICT architectures have had little problem in trying to allocate least privilege to users. Control system architectures are different in the sense that, historically, single users would have authoritative access across the entire infrastructure and access was limited by physical countermeasures. When a system does not have an inherent access control capability, functionality to provide access control should be added somewhere in the communication chain. For example, network firewalls protect systems from inbound communications by preventing all communication methods except those required for the desired functionality.

The following list describes the basic security features required from any system to protect access to critical resources, services, data, or communications. Although reused from the good practices associated with ICT, the principles can be extended to the control system domain. Every system in the chain of communication between a user and the end goal should possess all these capabilities. Although these capabilities can be deployed in ICT architectures with little or no complexity, the requirements associated with availability and integrity (often combined with non-standard technologies) demand that deployment in control system domains could be nontrivial.

- *Identity establishment* - a framework for providing or exchanging unique identities. User IDs are a common method, but are difficult to make both unique and meaningful in large environments and are not portable. E-mail addresses are portable because they incorporate the organisation name, but are difficult to make both unique and meaningful in large environments. Certificates are the preferred method. They can be federated (are portable) and the identity is an amalgam of many bits of identity data, rather than a limited user ID string, allowing them to be both unique and meaningful for very large groups.
- *Identity validation* - a mechanism designed to ensure the identity of each party to the others and one that cannot easily be forged. This can be as simple as a user password (called a 'shared secret' in the case of systems), but passwords are easily guessed, divined, or stolen; and therefore, identities that rely on them are easily forged. The preferred method is a one-time-password token or a password protected private key associated with the subject's signed, public certificate.⁹
- *Duress alerting* (optional, but recommended where available) - an Identity Validation solution that provides secondary credentials to users who can supply them to an attacker when needed. The use of secondary credentials will provide access to systems, data and services, but will alert operators and authorities that the user is under duress so that they may take appropriate action. This allows the user to protect themselves from a physical attacker who demands credentials to gain access to critical services.
- *Roles* - groups of users or systems organised according to their responsibilities. This can be centralised or distributed. Modern systems are centralising this function more and more, though there are designs that can provide decentralised role management.
- *Access rules* - lists of access rules that govern which groups and individuals are allowed access to certain resources at what times. These should be carefully planned to ensure that they follow the principle of least privilege. This function can be centralised or distributed. Modern systems are centralising this function more and more, though designs can provide decentralised access rule management.

Logging and reporting

This practice document, thus far, has focused on preventative security measures. Yet this is but one element of secure operations. Critical to any security architecture is the detection and reporting of anomalous activities. Preventative measures usually cannot be exhaustively implemented and certain environments present difficulties for some of the techniques mentioned above. Even though control system domains have been designed to detect anomalous activity (usually by events and alarm management), the incorporation of anomaly detection for users and services has traditionally been the responsibility of IT and not the automation group. For example, an organisation may not have the resources to deploy a completely federated identity

g. 'Strong authentication' is a poorly defined term and generally means authentication that is better than passwords. Authentication can be done in several ways that are independent of one another. Each method uses a different property of the user - something they know, such as a password; something they have, like a token or certificate; and something they are, as in some form of biometric, like a fingerprint or retina print. Adding a second password to secure access is the same as one password that is as long as the two combined and requires only a marginal increase in effort in order to crack it. However, using multiple factors, such as a password protected certificate, or a retina scan and PIN, requires an attacker to successfully attack two separate mechanisms which are dissimilar. For a more thorough treatment of authentication mechanisms and their relative strength, please refer to the DHS Practice document entitled *Improving Industrial Control Systems Cyber Security with Defense-in-Depth Strategies*.

management solution and ensure that it can work with all the legacy systems and applications that are already in place. Sometimes, patch management cannot keep a system up to date because business critical software will be adversely impacted by an update and thus security vulnerabilities may remain in place.^h Also, network encryption cannot always be deployed on WAN links because it introduces too much delay in the communication, which prevents some real-time systems from operating correctly. In some critical infrastructure control system operations, real-time or near real-time data acquisition is a mandatory requirement related to availability and system sustainability.ⁱ

To ensure that preventative measures are not bypassed, or to identify when someone or something gains access to what they shouldn't because of a deployed preventative control, it is necessary to monitor activity on or directed toward important assets and services, analyse them for specific content and alert someone when an anomaly is detected. The following list shows some of the types of events that should be identified by your logging system and for which an alert should be made to someone or something that can appropriately respond.

- Monitor failed authentication attempts - all devices or processes that require identity authentication should log and/or alert when an identity validation attempt fails.
- Monitor successful authentication attempts from different sources - If available, all devices or processes should log and/or alert when the same user logs in simultaneously from two different source locations.
- Monitor successful authentication under duress - for critical systems, consider deploying an authentication mechanism that supports duress codes. This allows a user under duress to log into a system using a secondary credential, but alerts that the access was performed under duress.
- Monitor failed access attempts - all devices or processes that manage access control to communications, data, or services should log and/or alert when access is requested that is not allowed.
- Monitor successful access attempts - all devices or processes that manage access control to communications, data, or services should log when access is requested and allowed.

Procedures should be established for handling these alerts. With regard to modern control system architectures, the demand from the user community has forced vendors to create advanced monitoring capability in their solution technologies. Generally, the functionality is presented from the perspectives of performance monitoring rather than security. Any functionality that supports the understanding of failed attempts, improper authentication attempts, failed connection attempts or any other abnormality can be used to support cyber security risk mitigation. At least, the collection of failed authentication or failed access attempts can be used to support a robust remote access policy. As stated before, this approach can be used for both users and services within the control systems information architecture.


h. The issues related to patch management in control systems have been researched extensively and the reader is encouraged to review *Patch Management for Control Systems* [csrp.inl.gov/Documents/PatchManagementRecommendedPractice_Final.pdf](https://www.csrp.inl.gov/Documents/PatchManagementRecommendedPractice_Final.pdf)

i. Cryptographic solutions and their impact to industrial automation are well documented in *Cryptographic Protection of SCADA Communications. Part 1: Background, Policies and Test Plan* <http://www.aga.org/NR/rdonlyres/B797B50B-616B-46A4-9E0F-5DC877563A0F/0/0603AGAREPORT12.PDF>

Network architecture security

Again, remote access is a subset of network architecture and the techniques to secure it are the same techniques used for securing the whole network. After putting all the items mentioned above together, an example illustrates how to secure remote access into an ICT infrastructure. First, create a list of assumptions:

1. Multiple types of 'remote' users require access into the network for different reasons. Some require access to non-critical systems, others require access to multiple critical systems and others require access to only a single critical system.
2. Multiple WAN connections compose the corporate network, so not all critical systems are necessarily located in the same physical structure.
3. Attackers are prevented from physically accessing resources that are inside buildings or compounds. In these examples, all attacks are against remote access mechanisms as defined earlier.



(Figure 4: Example of poor network security)

As shown in figure 4, the network has poor network architecture. The security of this network is poor as the following security elements are not present:

- Segregation of critical systems;
- Standardised method of remote access;
- Visible distinction between corporate roles to ensure access rights can be granular;
- Acknowledgment of the risks associated with unprotected telecommunications that traverse unprotected physical space.

Several examples of remote access are shown in figure 4:

- Vendors connect to individual systems through direct dial-up modem connections to provide support to critical business servers;
- Satellite offices connect to the main site through telecommunications WAN links that are not in the physical control of the organisation;
- Both modems and VPN tunnels are available to corporate users and vendors;
- Wireless access is available to different user groups.

Each of these access methods is susceptible to attack.

Improving security

Using the principles and techniques described in the section 'Applying good practices', the security of this implementation can be significantly improved by tackling each of the vulnerabilities one at a time.

Identify business critical items - This is the single most important part of all security plans, whether for network security planning, safety, secure application development, or reconfiguring systems to improve their local security posture. Knowing how to classify systems in alignment with the probable business impact if they fail allows one to group them physically, logically and conceptually to address their security issues in a coherent manner.

Organise network architecture - Create zones where systems, data sources and computing services are placed in groups whose members share similar security requirements. In the example below, web services have been moved to a demilitarized zone (DMZ) to protect them from both internal and external attackers. Critical services have been segregated to a secured zone to reduce susceptibility to attack from the corporate network should an attacker infiltrate the network that far.

Implement strong and granular access controls - This must be done on business critical systems and services as well as on network perimeters surrounding groups of systems with similar security requirements. The principle of least privilege must be followed when constructing these access controls. Granular access controls should be considered on every system, every application and every network border protection device such as firewalls and VPN concentrators. Figure 5 shows how communications into and out of the 'Business Servers' zone can be heavily restricted. For example, users might be granted broad access to the application servers' web interface, but administrative access could be configured such that

administrators would have to first connect to a special administrative server using a cryptographically secured connection and then manage the Business Servers from that one 'jump' server.


Standardise on one method of single-user remote access - The recommended method is to use a modern VPN, because all modern implementations of VPN technology include strong user and device authentication mechanisms. They also provide an extra layer of network access control, allowing different users to tunnel through to different systems using only allowed communications protocols. In Figure 5, all modems have been removed and vendor support access has been moved to VPN technology. This has normalised the security protocols for all external to internal access and reduced the number of possible points of attack.

Enable remote access only when required - Normal industrial control systems users may require on-demand remote access, but vendor support may only require remote access rarely. One can disable vendor user IDs until they are required to be enabled and then disable them once again when they have completed their task. This technique can be applied to any untrusted group of users who require only intermittent access to corporate resources.

Use strong authentication credentials – Users needing access to critical systems and services should be required to use strong authentication methods, such as one-time password tokens or certificates. These are also known as multifactor authentication methods. Devices such as wireless access points and microwave towers should use strong authentication to identify themselves to each other, so that access rules can prevent unauthorised wireless devices from participating in communications.

Protect wireless communications - Use modern cryptographic techniques to protect WLAN and microwave communications. This makes the wireless communications methods much harder to attack. Injection and replay are very difficult and eavesdropping and hijacking become almost impossible. Most microwave implementations and all WLANs have advanced cryptographic capabilities. The communication lightning bolts in the diagram have been changed from yellow to green to indicate that the end points use strong authentication to validate their identities and use encryption to scramble communications.

Logging, monitoring and alerting - Critical devices should log activity and alert on anomalous events. Devices in this category include all servers, applications, communications equipment and network perimeter devices, which are critical to business operations. In addition, the deployment of monitoring devices captures and examines network traffic for anomalous behaviour. They are typically deployed behind network access control devices to ensure that changes to communications behaviour do not go undetected.



(Figure 5: Security countermeasures)

Remote access in control systems architectures

With the growing interconnectivity between control systems architectures, corporate architectures, peer sites and other operational entities, organisations have had to abandon the traditional (and sometimes ideal) concept of total domain isolation. Realistically, industrial control systems have always had some aspect of remote access play a part in operations. As discussed before, vendors have had access to support their systems and the communications infrastructure was traditionally quite extensive so that it supported data control and acquisition from long distances. The mechanisms for data acquisition involves several different types of communications media, many of which were not dedicated to a single utility but were shared among some number of different entities. A number of the security functionality and concepts that ICT has used can be leveraged in control system architectures. The challenge is how to apply cyber security good practices to remote access programs such that the solution supports the requirements for business operations.

This section is designed to examine remote access solutions in industrial control systems environments and show how the lessons learned in ICT environments can apply here. Attention should also be applied to examine the differences between the environments and how they constrain the choices of security techniques available.

One of the best ways to understand the differences in requirements for remote access between the business ICT domain and the industrial control system domain is to look at a comparative matrix. These differences are the foundation to address some of the more unique elements that impact strategies for deploying remote access and control system architectures. The following table shows the comparison of key remote access requirements and the differences between ICT and control systems domains.

Remote Access Requirements	Information & Communication Technology	Industrial Control Systems
Direct internet access	Not common, protected by firewall and intrusion prevention system.	Control system components are sometimes exposed directly to the internet
Modem access	Rare	Common/widely used, legacy equipment
Leased lines	Common/widely used	Common/widely used
Remote access through Firewall with virtual local area network (VLAN) segmentation	Common/widely used	May be segmented from corporate network, little segmentation within control system network
Remote authentication	Multiple form factor	Usually single factor, some multifactor solutions ineffective due to speed or process control reliability requirements
Remote authorisation	Usually role based or based on confidentiality of data	Usually a single level with full authorisation
Access control lists	Usually well defined	Not always utilised or clearly defined. Not maintained as it should be
Audit trail	Robust and actively monitored	Sometimes limited capability, not actively monitored
Authorisation server	Maintained and monitored in a DMZ	Authorisation servers deployed directly on the control system network
Demilitarized zone	Access through DMZ only	DMZ not always deployed
SSL Encryption	Commonly deployed	Rare
IPSec	Commonly deployed	Very rare

(Table 2: Comparison of control systems and ICT remote access requirements)

Remote access security considerations unique to control systems

Culture has always played a part in how cyber security is implemented in control systems environments. When security foundations are based in the complete and total isolation from untrusted domains, the migration toward creating security solutions that account for interoperability can be challenging. Regardless, the basic attributes associated with control systems functionality, attributes that are based on requirements for high availability and data integrity, create opportunities to leverage proven security technologies and adjust for operational requirements. The perspective of those responsible for creating remote access solutions that allow for direct connectivity into industrial control system operations may not perceive cyber security as a critical concern but rather how the access solution can be managed to maintain critical operations. When trying to address the security component, many direct connections are justified by the perceived obscurity of the system and the risk is mitigated by

the assumption that little understanding exists of how the system actually works. As has been repeatedly proven, this approach is not only dangerous but can lead to some significant operational risk (i.e. the lack of awareness that mission-critical systems are directly connected to the internet).

A vast majority of control systems environments are deployed in domains that are considered to be critical infrastructure. Risks to these environments are not limited to the company operating the infrastructure. Remote access to a control system does expose some aspects of the architecture to remote manipulation. Remote access may be an exploitable attack vector that adds extra risk regarding the availability of the control system. The introduction of security for remote access cannot impede or degrade the normal operational processes that are critical for the control system to function normally.

For example, the remote access security implementation will have to consider the necessity for real time operations. Surprisingly, many organisations fail to recognise the realistic impact security can have on real-time operations and will often discount the possibility of deploying security countermeasures without appropriately analysing the impact performance. Many control systems environments need to operate in real time, with some environments requiring sub-millisecond polling. Any latency that is created due to the deployment of a countermeasure, such as encryption, may negatively affect the overall process and cause unnecessary delays or shutdowns. As well, much of the data on a control system can be deemed non-confidential and thus the lack of need for encryption within the communication between critical system components.

Many control systems environments are geographically dispersed and may even cross international boundaries. The nature of these deployments requires that many field locations are unmanned and the requirements for availability often make the remote access solution address connectivity more than security. This may limit some procedural-based security protections such as allowing only temporary access to the system. Adding security functionally that may slow down the management of the field equipment, such as calling a help desk to enable remote access, may be justification for keeping an 'always on' remote connection.

Because many control systems environments have the requirement to operate in real time, the demand to quickly connect to a system when necessary is crucial. Often, operators may feel impeded by the multiple steps required for remote access and will either want to remove some security features that slow down their connection process or create workarounds to expedite connectivity. A good example of this includes an organisation maintaining the use of the default administrative credentials involved in remote access, or the creation of passwords that are not complex and are not forced to change on a regular basis. As stated before, when operators are working under duress and system survivability is paramount, many users will want to (or be required to) connect to a remote system as quickly as possible and do not want to have to worry about connecting with a complex password or one that they have not used for many years. These practices, in addition to those that involve using the same password for every field device, greatly reduce the chances of a remote access capability being secure. The need for remote access to be as quick as possible can justify not adhering to added robust security protections.

Adding to the complexity of security requirements for control system architectures is because not all remote connectivity will take place in or from a controlled location but rather from an area that can be impacted by environmental effects. This can create a problem when attempting to use some advanced features for secure remote access, including ones that require operators to use multifactor authentication. Many control environments are deployed in factories, large processing plants, warehouses, isolated stations and even far outside of a six-walled building.

The environmental issues associated with working in these locations can create situations where an operator is unable to take advantage of advanced authentication technology. Many organisations face the challenge in trying to deploy remote access security strategies that are commensurate with the value of the system being administered. However, the usage of biometrics or some other technology that requires a clean environment is useless, or operators may be unable to be free of dirt, grease, oil, or some other element that prohibits the effective use of the technology. For example, two-factor authentication using advanced biometric thumbprint scans can be affected by the dirt or oil and grease that build up in these types of environments. Taking advantage of voiceprint technology can be very difficult, because some of the background and residual noise of working in a large processing plant or outdoors in an open air environment may affect the voice wave length for recognition.

Finally, from a pure technology perspective, many control system devices or implementations may not have the capability to effectively use even basic security features such as authentication or authorisation. With control systems having a larger than average life cycle, some upward of 20 years, the incorporation of effective remote access security countermeasures is just simply not plausible. Even though the asset owner may make repeated requests to the vendor, the cost associated with implementing the effective remote access solution is larger than the purchasing of a new system itself. This situation does not favour the operator but rather the vendor, leaving the operator little choice in creating aftermarket solutions. Secure remote access can be accomplished, but asset owners can expect to require the full support of the vendor to help secure existing remote access capabilities.


Applying good practices

Current interest in secure remote access has created a situation where there are a multitude of options in defining what a good practice should be. Although specific good practices are associated with highly regulated ICT systems, such as those used in the federal government, financial or health care sectors and these practices are not straightforward. As discussed earlier in this document, there is tremendous complexity in trying to create a remote access good practice that is going to be applicable to every possible control system architecture. Asset owners are in a position, however, to create effective baselines from which secure remote access programs can be derived. By understanding the accessibility requirements in their organisation and cross correlating with both the users and services required for remote connectivity, operational mandates can be met in a manner that accommodates for the appropriate protection of critical information assets.

Most operational environments have limited choices for the technology that can be used for remote connectivity. When combined with the recognition of the critical assets needing to be accessed, the task of defining guidelines can be straightforward. Guidelines that are immediately appropriate to control systems environments include:

- Undertake a formal threat and risk assessment;
- Eliminate all direct connections to critical operational assets;
- Secure modem access beyond default means;
- Use DMZs to segregate business and control architectures;
- Establish user-specific authentication servers;

- Create a security assurance policy for all remote access;
- Use only full tunnelling cryptographic technology;
- Use a password policy specific to remote access elements;
- Wherever possible, use multifactor authentication;
- Use role-based authorisation levels;
- Use dedicated hardware and software to support the remote access solution.



(Figure 6: Basic industrial control systems network with no security)

Figure 6 shows a basic industrial control systems network with no security. The whole network, including corporate, is protected from the internet, but there are no security mechanisms in place to protect the industrial control systems environment from the corporate computing environment, business partners and remote office locations. The next several sections will illustrate the different controls in the environment and will conclude with reasonable security controls in place.

Periodically undertake formal threat and risk assessments

Organisations need to understand their threats and risks. Risk management is a common element of project management and is used by organisational managers everywhere. However, project managers are concerned about on-time delivery and managers are primarily concerned about margins or efficiencies and the risks routinely considered by people in these roles reflect their perceived priorities. The topic seldom turns to risks that emerge from threats of system failure or deliberate sabotage because of absent security controls. As a result, these topics


typically receive nominal treatment from managers who then do not acquire a deep understanding of the cyber threats to their industrial control systems infrastructure.

In order to solidify an organisation's understanding of these risks so that managers can make risk-aware project or operational decisions, formal Threat and Risk Assessments (TRAs) should be conducted on industrial control systems environments on a regular basis.

Eliminate all direct connections

Because of the security risk and ease of attack path, the elimination of direct connections to control system architectures is strongly recommended. Organisations need to consider that the prevalence of flat networks control system architectures could create unrecognised direct connections. As many control system architectures trust many operational resources via remote access channels for functionality, the control system could be compromised quite easily by exploiting direct connections. Field devices and equipment require special attention and the access requirements associated with these assets make the removal of direct connections difficult. In the event that organisations cannot eliminate direct connections, countermeasures that involve authentication and authorisation should be deployed. At least, wherever possible, authentication and authorisation mechanisms on the field equipment should be enabled; and if there is no requirement for remote access to the field equipment, such capabilities should be disabled.

In the rapid expansion of control systems environments, some organisations are left with little choice for remote connectivity. Devices, such as relays, intelligent electronic devices, remote telemetry/terminal units (RTUs) and Programmable Logic Controllers (PLCs) should not be connected to the internet even if they have a username and password set. If these devices do require internet connectivity, organisations should proactively ascertain what functionality can be removed to lower the risk profile and work with vendors to establish unique access capabilities that significantly limit the exposure of the control system. Entities should leverage procurement language guidance to ensure that security capabilities are inherent in the technology being procured and that vendor assistance is available to support secure remote access to the device. When direct connections are required to critical information assets, a good practice is to remove functionality that is not required for operation but, if leveraged by an adversary, would result in a security incident.



(Figure 7: Eliminate 'direct connections' by adding firewalls around the industrial control systems environment.)

Figure 7 shows an industrial control systems network that is partially isolated from all other networks. This is similar to the old architecture, where the industrial control systems network was completely isolated from other environments. However, the firewalls can be configured to pass appropriate traffic into and out of the industrial control systems zone.

The modem pool, if it continues to be required, is moved outside the industrial control systems to a DMZ-like network to prevent the modems from compromising the industrial control systems environment.


Secure modem access

If modems are not used, they should be either removed or disabled from the device or system. If a modem is deemed necessary despite the available alternatives, they should not be left on and should be deployed using recommended practice guidance.³ For further guidance on securing modems, please refer to the DHS practice guide entitled *Securing Controls System Modems*.

Create a physical and logical DMZ separating corporate and IT environments

All authentication servers and access servers should be placed in a DMZ. This separate logical and physical network prevents direct access into the control domain from the corporate or external environments. Although it does not impede business functionality, it can greatly increase the work effort of an attacker. By locating the authentication servers that must be used to gain access to mission-critical assets, additional levels of security can be implemented. These authentication servers should be located on a completely separate VLAN to support the

isolation of these networks. When this approach is used in tandem with the isolation of critical operational assets, such as field devices, specific guidance that provisions in the credentials for remote access can be made easier. Much work has been accomplished as it relates to creating DMZs within control systems architectures. For further guidance on securing network architecture, please refer to the DHS practice guide entitled *Improving Industrial Control Systems Cyber Security with Defence-in-Depth Strategies*.



(Figure 8: Add a control systems' DMZ and segregate field devices from industrial control systems network.)

Figure 8 illustrates network segregation. The network is divided into functional areas and the architecture is similar to that of an internet connection from corporate.

Create separate authentication servers for separate roles (vendors/integrators)

Asset owners should recognise that operational assets that contribute to control system functionality can be considered trusted or untrusted. Operators and engineers have traditionally been considered trusted elements and when inside the operational domain of the control system, their access is restricted based on policy. External elements, such as vendors and integrators, create a situation where the electronic boundary is not easily perceived and may be considered untrusted. Although the trusted operators (internal) and untrusted support entities (external) may share the same architecture once connected, the mechanisms for provisioning the access into the environment should be different. Not only will this help with security auditing and incident management, but it will also provide a much more granular capability to restrict access based on user profiles.

A separate server should be created for authorised users that come from an external organisation such as vendors or integrators. This creates the opportunity to create vendor-specific access levels and they can provide for control mechanisms that limit a number of

different factors that range from time of day to traffic patterns. Dedicated authentication servers for different roles allow for swift mitigation and security countermeasures in the event of a cyber incident. This will also give the asset owner the ability to dynamically lock down the vendor or integrator access that can be quickly granted back when needed or on a case-by-case basis.

Depending on the architecture that facilitates access into the control system environment servers can be deployed at a number of different locations connected to the external communications infrastructure. Furthermore, the authentication servers can be deployed to support existing or future media formats and can be deployed such that the external resource must comply based on agreements set up during system procurement. As standard good practice would indicate, the authentication servers should have different user names and passwords than the user names created on the corporate or ICT network.

Enforce a security assurance policy for all remote access

All users of any remote access system should sign a policy statement that supports the expectations related to the provisioning of remote access. This policy will clearly define legal and acceptable use conditions associated with use of the system to which the remote access is facilitating a connection. The security assurance policy should be formalised, documented and approved by a senior manager and should have the flexibility to be customised to accommodate different users and different requirements.

A good practice is the development life cycle of the remote access policy, which aligns with existing ICT security policies that are mandated by overarching business functions. Modifying the security policy to accommodate control systems architectures helps define the scope under which remote access users must function and clearly states the organisation has an understanding of the consequence associated with the misuse of the remote access channel.

Specific to the remote access policy are detailed definitions on device management and how devices are identified as essential to core operations of the control system architecture. Unless specifically defined, access to critical devices is to be read-only operations and deviations from that standard are restricted to other categories of operators. If the remote access policy and threat model assesses the essential need for remote command, control and change access, then appropriate mitigation techniques should be identified to ensure core operational devices are able to be effectively monitored, alerted and archived. The policy should also address the purpose, scope, roles, responsibility and management of all remote access requirements. This must be completed before any access is granted.


Full tunnels

Securing remote access communications to ensure availability and integrity demand security countermeasures that prevent an adversary from getting access to the data. In the control system space, much of the critical data are in real time and a compromise of the command and control instructional information could have significant impact on the system. This fact suggests that whenever possible the remote access that supports the operation of the control system should be secured and authenticated by cryptographic means. Although the mathematics involved with contemporary cryptographic solutions is beyond the scope of this paper, of all the methods for point-to-point encryption that can be used, the best solution in the control system arena involves 'full tunnels.'

Organisations looking to deploy cryptography have many options at their disposal but the preferred solution doesn't always preclude an adversary from injecting himself into what is

assumed to be a secure connection. The use of full-time roles prevents this type of attack and prohibits the attacker from acting as a gateway between the control system network and the network from which the trusted resource is coming. In addition to ensuring that both authentication and authorisation occur, the solution can reduce the risk relating to the injection of viruses and worms into the data stream.

Most entities which use full tunnels create a remote access solution where initial authentication and authorisation must be to the corporate network first, then to the control system DMZ authorisation server. This provides some protection against a compromised corporate account, thereby thwarting one of the more popular methods of control system compromise seen today. As adversaries continue to exploit the relationship between business and control systems environments, any remote access solution that prohibits the exploitation of corporate trust while maintaining critical access to control operations is a valuable one.



(Figure 9: VPN Remote Access)

Figure 9 illustrates the use of encrypted tunnels, or VPN technology. In some cases, as with modems, it can replace the access method with a more secure mechanism. In other cases, as with network-to-network communications, it can add needed security by encrypting traffic or requiring endpoint authentication.

With regard to cryptographic solutions, the discussion is not complete without paying attention to the issues surrounding latency. Indeed, almost every cryptographic solution available that is designed to secure point-to-point communications has the potential to impact data throughput

on a network. This latency is a function of not only the type of cryptography being used but the algorithm being used to secure the data. Point-to-point cryptography uses symmetric key algorithms where the key that is used to encrypt the information is the same key used to decrypt the information. By its very nature, symmetric key cryptography is not as secure as asymmetric key cryptography, but the attributes associated with symmetric key cryptography as related to minimal key management is attractive for large-scale environments using numerous disparate devices. For control systems operations, these attributes make a symmetric cryptographic solution and the best choice, even though the compromise of the single symmetric key could render the solution useless.^j Even with this risk in mind, as well as the possibility to increase latency up to 30%, the symmetric solution remains very popular with many vendors deploying the Advanced Encryption Standard (AES) in either the 128 or 256-bit version.

With regard to remote access, organisations need to recognise that the usage of the algorithm does not necessarily have to be a limiting factor in choosing a successful solution. In some cases, advanced developments in compression algorithms used in cryptographic solutions can actually reduce the latency and increased speed of communications beyond what they were prior to encryption. Numerous vendors actually provide technologies that were developed to support field technician remote access into mission critical devices and have extended their solution to create secure channels for remote data acquisition from centralised servers.

The one caveat associated with these unique solutions is that the algorithms used may not necessarily conform to recognised standards and some of the cryptographic solutions may not have been fully vetted by the research community. Still, for some entities, the comfort of having some data security without impacting the rate at which vital operational data can be acquired is more than enough to justify its implementation.

Password policy

In an ideal deployment, authentication mechanisms should be chosen based on the criticality of the system being accessed and should include not only passwords, but other mechanisms as well (see the section on 'Two form factor authentication'). When using passwords, a secure remote access system should enforce complex passwords of 8 to 25 characters that are a mixture of upper and lower case letters, numbers and symbols.^k In addition, corporate policy should demand that these passwords be changed at a rate that is commensurate with the value of the system being protected and regular audits of the strength of these passwords should be done as part of the organisational cyber security program. The corporate cyber security policy should dictate that these passwords are never shared and that each user ID is unique across the entire system.

For remote access programs, the creation of user ID and associated passwords should be done by a process that is unique to the type of user or service and their associated enclave. Organisations are encouraged to classify users and services into levels of trust and deploy a password policy based on the risk level of the user or service.

j. Many security solutions in industrial automation use both symmetric and asymmetric solutions. However, some standards mandate that only symmetric solutions can be used (i.e. ANSI C12.19 and C12.22 for smart metering communications). 256-bit AES encryption should be implemented from client to corporate DMZ to control system DMZ. Encryption levels into the control system domain should be tested and verified and will be dependent on the capability of the end control system devices to support 256-bit AES encryption.

k. This guideline is a recommended best practice for general ICT security and the authors recognise that the creation and use of a 25-character complex password (although ideal) for day-to-day human machine interface operations may be inappropriate. The recollection and usage of such a password under duress may create circumstances that are unacceptable from a safety perspective.

The password policy needs to extend to the remote access technologies that are provided by the vendors. The passwords must also be applicable to control system technology that may provide for remote connectivity. Stated another way, if the vendor provides remote support technology, such as modems, as well as manufactures embedded remote access functionality in ICT field device technology (i.e. RTU, PLCs), the organisation should have a password policy for the vendor's remote support access and one that defines the provisioning of username and passwords for field equipment. Traditionally, the flexibility of vendor password policies has been limited and changes to their standard procedures have only been considered when requested by very large customers. Many vendors do not provide the capability for any customer to change administrative credentials required to access field equipment. Moreover, some vendors still hardcode supervisory passwords into their control system making it impossible for a customer to change them. These issues, when combined with the historical cultural barriers associated with cyber security and control systems, create situations that elevate the cyber risk an organisation can have.

Organisations developing mature cyber security programs to mitigate the risk associated with remote access should work with their vendors. As part of the procurement process, asset owners should demand that the capability to update administrative credentials for control system equipment be embedded in the vendor solution.

Two form factor authentications

Authentication is the process of validating identity. Methods of validation can be broken down into three categories, or factors. They are:

- Something you know (a secret, like a password or PIN);
- Something you have (a token or object that is unique);
- Something you are (biometrics, such as fingerprints, retinas, or gait).

Attempting to improve security controls by using two passwords to access something is similar to using one password that is as long as the two combined and the method for 'cracking' one is the same as for the other. However, using multiple authentication mechanisms from different categories, or multifactor, means that the different mechanisms are unrelated to each other and require independent means of 'cracking' them.

Critical control system assets, especially those that are only accessible through remote means, should be deployed with dual-factor authentication to protect the link or at least protect access to the information resource or device. Today, for remote communications between main control centres and backup facilities to occur across a trusted channel, access into critical operations is accomplished by multifactor authentication. This process is also very common in situations where field technicians are connecting remotely to critical operational resources to support system continuity, restoration, diagnostics and upgrades. Clearly, the compromise of these access channels by an attacker could result in some significant damage, because once access has been granted, instruction sets will be interpreted as coming from trusted sources.

The use of centralised servers for multifactor authentication is critical to the success in securing remote access. The software, which is loaded on information resources, needs permanent servers or mobile tablet computers and should be standardised across all remote connections. This standardisation greatly empowers an organisation to create access control lists for remote access and helps provision the appropriate audit functions that ensure only authorised activity is happening in the operational domain.

Although it is inappropriate to make the assumption that every organisation needs to prevent attackers from capturing static username and passwords, which can be reused by the attacker at a later time, organisations must perform appropriate risk assessments to define the level of protection required to secure the credentials involved in remote access. When dealing with primary critical systems, some organisations prefer to use multifactor authentication over cryptographic channels and force an initial authentication connection to happen on demilitarized servers. This is just one example of how defence-in-depth strategies can be used to secure remote access while not impeding business requirements.

With the deployment of security counter-measures, particularly those that are deployed to administer access control over operational assets, an associated cost of ownership goes with managing the supporting technology. If a multifactor authentication mechanism is to be used, then good practice dictates that a standardised method be available for administering remote access and that could entail a dedicated server or services. The location of these servers will dictate who incurs the costs associated with the ownership and management of this service. Like managing firewalls and intrusion detection capabilities, predefined responsibility of the servers will help expedite a more robust remote access solution. Failure to do so could result in improperly managed access control capabilities that could be exploited by an attacker.

Authorisation levels

Depending on the outcome of the threat and risk assessment, the secure remote access system may require role-based authorisation levels. Read-only access is the default user level and should be the de facto user level in the same way that the access control lists in a newly deployed firewall prohibit access of any kind (until access control lists have been established to provide only authorised connections). Access to the full network is limited and the concept of least privileges is used. Users must escalate their privileges to make authoritative changes on the network, control system devices, or any critical function.

Organisations should deploy remote access capabilities in a manner that guarantees that the access to a single point will not automatically provide access to multiple points in a control system environment. Although apparently contradictory to common sense, organisations need to recognise that the rapidly emerging interoperability with historically untrusted networks demands that attention be paid to the consequences associated with the exploitation of trust. Organisations that simply provide a single level of authorisation for remote access users, especially ones that need to support operations across different enclaves, risk an unseen degradation in security posture.

By the simple allocation of identity and value to operational assets, organisations can reuse remote access authentication and credentialing capabilities for any operational domain or asset. This would support the policy directive suggesting the secure remote access system does not allow access to multiple control networks without re-authentication and authorisation. Where possible, identity management techniques should be deployed to separate the identity of the same user's corporate account from the control system account. Changes to all aspects for authorisation levels and identity data information are auditable and archived for managing change, incident response and forensics. This approach has been proven to work in several sectors and in several mission-critical domains.

Dedicated client hardware and software

One of the fundamental elements of any remote access solution is that the users of the remote access system will be empowered with both the software and hardware required to connect. Yet threat and risk assessments show that both the hardware and software of a dedicated remote access client is the target of an attacker. Once an attacker has access to client hardware or software he or she is only limited by the ability to get past the local security of the system. Traditional approaches that provide remote access software on the standard issue mobile computing device have enabled attackers to not only access operational data, but to access operational networks as well. Recommended practices suggest that the technology for doing the business and the technology for getting access to the business should be separate.

As part of the remote access solution, the organisation should provide a dedicated personal computer (PC) or laptop for VPN access, one that is centrally managed and hardened with a baseline image specific for remote access needs. The PC or laptop will have appropriate cyber security counter-measures, such as antivirus or host-based intrusion detection systems software and can even include security counter-measures that have been designed to meet the unique requirements for field operations and control system functionality. The PC or laptop policy and procedure should be clearly spelled out that the use of this system is for secure remote access. Another consideration would be to use the MAC address of the system as an additional authorisation factor. Many organisations have found it helpful to designate remote access resources to be tuned to specific hardware identification metrics that are physically attached devices.

If a PC or laptop cannot be provided, then the investment to provide antivirus and spyware removal tools for a user's home computer is suggested. If providing dedicated ICT resources to remote users is not a viable option, supporting the configuration of personal computers to accommodate the specific security needs for remote access should be performed. Organisations which require remote access to control systems operations should be able to have any computer adjusted to a baseline security standard as well as deploy cryptographic and disk encryption security measures to facilitate the protection of operational data. While never recommended, if the user's home computer is going to be used for remote access to control operations, then secondary and tertiary identification/authorisation methods, as well as up-to-date antivirus, anti-Trojan and anti-spyware software measures, should be established to ensure security. Regardless of whether the organisation is going to use dedicated resources or personal home computers, specific education and outreach pertaining to access procedures and security safeguards must be taught. Without appropriate cyber security education, covering threats, risk, vulnerabilities, consequence and the appropriate use of remote communication technology, the security element of the remote access solution is incomplete.

Session termination

The concept of establishing a session is fundamental when discussing remote access. The management and protection of that session is fundamental when discussing secure remote access. No discussion of secure remote access is complete without brief mention of session and session termination. Although it is conceivable that some remote access sessions will be established in perpetuity, usually to facilitate communication needs related to availability, remote access solutions need to be deployed with viable and effective mechanisms for terminating sessions. Session termination is a mandatory element of any secure remote access solution and organisations should be able to facilitate session termination either on request or automatically due to system configurations.

Some of the more serious oversights related to remote session termination that have been observed in control systems environments have led to the creation or the unintentional support of covert channels and situations of denial of service. Regardless of the method used for remote connectivity, communications channels need to be configured to terminate under certain conditions. These can include, but should not be limited to:

- Termination of the session after a set period of time;
- Termination of the session upon reaching predefined triggers (time of day, day of week, on request, data type, file size, etc.);
- Termination of the session based on inactivity;
- Termination of initial response capability based on number of failed attempts;
- Termination of session based on cryptographic key exchange failures;
- Termination of session based on quality service;
- Termination of session based on confirmed tampering.

Managing remote access in control systems environments

When creating and deploying a secure remote access solution, an organisation will need to involve both ICT security and infrastructure teams along with control systems engineers and operators. In order to ensure security is embedded as early as possible, this cross domain cooperation is best started at the design phase of the secure remote access solution.¹ Figure 10 provides a notional diagram that illustrates secured remote access.


The purpose, roles and responsibilities should be clearly defined from the outset of the design of the secure remote access solution and should be reviewed, at a minimum, yearly. The organisation should clearly define their unique requirements and capabilities of their infrastructure and resources. Operational considerations should define and disseminate areas of responsibility that are only for ICT personnel managing the business and external access networks, only for control systems personnel and those responsibilities that are shared between them. ICT personnel should be responsible for the communication and security configuration of all access points and devices from the remote client to the control system DMZ. All remote access and remote communication from the control system DMZ into the control system network should be the sole responsibility of control system personnel. The servers, network segment and firewall that separate the control domain from the corporate domain and make up the control DMZ should be a shared responsibility with the control system and the ICT personnel. Proper training and understanding of both domains shall be established for personnel managing the shared responsibility domain.

Security management of the remote access solution should follow defence-in-depth techniques with multiple layers of defence and should incorporate best-of-breed elements from standards. Although not an exhaustive list, examples of layered defence techniques in relation to remote access are the following:

- Intrusion prevention solutions should be deployed on both the business ICT DMZ and the control system DMZ, leveraging operation-specific signatures and triggers defined by the remote access solution parameters.
- Provisioning of identities, authorisation and authentication should be a separate team that adheres to strict change management policies and procedures.
- VLAN or physical segmentation within and between the business ICT networks and control system networks, connected only through a robust firewall configured according to the principle of least privilege.
- Centralised log management and 24/7 monitoring of security events and logs for proactive incident response and more accurate forensics.
- Regular reviews and assessments of technologies deployed and policy and procedure enforcement.
- Patch management strategies for all ICT devices that make up the secure remote access solution, including remote clients, untrusted servers and access gateways and terminal services.

¹ Some organisations have found it very useful to include an option for the creation of a remote access solution into the procurement process, suggesting that vendors will create the solution as part of a value added.

- The use of anti-spyware, anti-malware and anti-virus services that are frequently updated on a scheduled basis.



(Figure 10: Secured remote access)

Case study: Marstad^m municipal operations

To better understand how remote access solutions can have a positive contribution on maintaining critical infrastructure control systems, it is useful to discuss a simplified use case. Although several different sectors could be used as the basis for an informative discussion on remote access, only three or four different sectors would require a solution that could address almost all the issues discussed in this practice guide. Although all the issues discussed in this practice guide deserve attention, not every organisation will need to consider each one of them. Here, a simple use case involving a municipal water treatment centre has been developed. This simple case study and the remote access elements within it can be extrapolated to any other sector as required. In addition, some of the issues related to the demands on a control system architecture regarding meeting growing infrastructure requirements are presented as well.

This case study is presented from the perspective of security risks, remote access requirements and solutions based on risk assessments.

m. Marstad is a fictional town.

Background

A municipal water treatment centre in Marstad has a control system architecture that manages several critical areas of the water supply, including source, treatment centre and distribution. The Marstad municipality is responsible for a large geographic area with a population of approximately 50,000 people. The municipality and surrounding areas are rapidly growing and new housing developments and businesses are rapidly increasing demand on the existing water supply. The Marstad municipality is expected to grow to just over 150,000 people in the next ten years. The water source for Marstad is located past the Marstad municipal boundaries that the actual treatment centre is in and crosses into two other cities contiguous to Marstad. The Marstad water treatment centre is within its own city limits and it runs the water pipeline to a storage facility in addition to remote customer sites.

The head office for municipal operations (corporate) and government works is located two blocks away from the Marstad City Hall building. To effectively manage the water system and the water treatment centre, both which serve very large areas, the municipality uses several remote access connections to facilitate water operations, vendor support access and integrator maintenance.

Control system architecture

As with all municipal works companies, the Marstad fresh water system's annual budget is partially dependent on factors outside the control of its management staff. As a result, the control system architecture is a mixture of design techniques that have been joined together over decades as funds became available.

When first installed in 1973, the electronic control systems were not centrally managed. The water treatment facility used its own control room in house. Pump houses, including the water tower, used electromechanical components until 1985, when the first PLCs were put in place. At that time, the water treatment facility was upgraded to be completely automated and required only a skeleton crew to monitor the system and troubleshoot problems. Pump houses were fitted with telecommunication links back to the head office, typically leased lines or ISDN. The water treatment facility was connected to the head office by a 256-K ISDN link.

Budgets in the 1980s were quite full. The economy was extremely strong; so despite tax cuts, the municipal coffers were well padded. At this time, Marstad water was able to justify and acquire approval to include leak-detection capabilities on its water mains. A program was put in place to add leak detection at 17 strategic locations in the existing mains and to include strategic placement of leak detectors during the addition of new community mains and during main replacement. Leak detectors on new installations can be connected through any telecommunications media, but those installed on existing or replacement pipe cannot typically be connected using physical communication lines. As a result, detectors use a mixture of copper, fibre-optic and radio frequency telecommunication methods.

Until 1992 the town was small enough for its supply to be equilibrated and managed solely by the water tower. In 1987, due to highly fluctuating water levels in the tower, the town recognised the need for a reservoir and it was completed in 1992, despite threats to defund the already approved project due to the recession in the early 1990s. The reservoir was built with its own control room, but it was unmanned. With the belt-tightening of the recession, all monitoring of reservoir systems was centralised to the head office control centre, which then had primary control of the reservoir, pump houses and the water tower and also had monitoring capability for

the water treatment centre. However, the control room at the treatment centre has been maintained, such that if necessary it can be used as a backup control centre in the event of an emergency.


As part of the reservoir project, the water delivery control infrastructure was updated to 'modern' PLCs, with a centrally managed DCS in the head office. However, the water treatment centre became a victim of its own success. Long periods of operation without incident combined with the budget crunch from the recession caused the treatment centre to lose its entire permanent staff and monitoring was moved to the central office, along with the monitoring of the upgraded pump houses. Consequently, by 1994 all monitoring of fresh water delivery was central, problems were managed centrally and staff was dispatched to remote sites as needed. The only employees who regularly enter facilities do so to remove waste from and supply products to the water treatment centre.

As the cost of leased lines grew through the 1990s, reprovisioning of telecommunications was undertaken, primarily as part of cost recovery efforts. The connections from the water treatment centre, the reservoir and 11 of the pump houses have been moved to the local telephone company's MPLS cloud. The remaining four were provisioned with internet connections in 2001 and small business firewalls with VPN capabilities were installed. The VPNs are used to tunnel across the internet to the head office to allow remote monitoring by the DCS.

The internal network at Marstad Municipal Operations is flat. The control system and corporate systems all exist within one 'security zone.' They have implemented a corporate firewall with a DMZ for their public website and customer service web applications, but though the internal network has many segments, there are no network access controls within it.

Prior to the advent of VPN technology, modems were used to allow ICT and industrial control systems administration staff to access systems from home. This modem pool has been reduced in size since adoption of VPNs for employee use, but they have not been eliminated entirely. In addition, the support contract between the industrial control systems system integrator and Marstad Water Distribution requires that the integrator has remote access to the DCS. Thus modems are installed on the DCS, the Data Acquisition Server and the Historian database. A similar contract with IBM has resulted in a modem being installed on the mainframe used for customer management.

Because the water source is shared between three political regions, they have all entered into a sophisticated resource sharing agreement. Consequently, they must share usage information to supply the data for reconciling actual use and are interconnected through WAN links.



(Figure 11: Marstad water distribution network - prior to security)

In figure 11, the Field Locations represent the pump-houses, the water tower, the reservoir, the water treatment facility and the leak detectors. Some of them connect via WAN links, which can be MPLS or radio frequency. Others connect via site-to-site VPNs over DSL and the internet.

All the VPN-enabled sites use IP to communicate between the PLCs and the systems at the head office. Those sites connected through the WAN use a combination of IP and serial to communicate with head office systems.

To leverage infrastructure and lower costs, the Marstad Municipal Network has been provisioned for use by all city employees. This includes, but is not limited to, City Hall, Marstad Water Distribution, Marstad Wastewater Company, public libraries and the Marstad District School Board. The Marstad emergency services each have their own network.

Internal to Marstad's network environment, administrators make extensive use of Virtual Network Computing (VNC) to remotely manage systems, including industrial control systems that support it.

Risk factors

For the Marstad operations, multiple network connections to and from myriad field devices are contained in the distributed control system. The steady growth in population has increased

demand for the supply of water to the municipality and as such, the municipal corporate office (where control operations are located) is experiencing an increase in demand for remote access to the control system. The demand is coming primarily from internal engineering and the requests are to meet the demands of operational requirements that must be met by engineering and the occasional vendor support.

To manage the water operations control system remotely, multiple external connections (sourced at different locations) come into the control systems via several pathways. The numerous remote access requirements create increased risk for water operations, but the largest security risk is created by the primary vendor of the control system. To maintain operational support and to ensure compliance with service agreements promised to Marstad, direct access to the control system Master PLC farm must be provided.

The procedure to facilitate this connection is straightforward and is not unlike the mechanism used by many organisations requiring remote support for their industrial systems. The switch that the master PLC is connected to, one that is connected to the control LAN, is simply connected to a router that has internet access. When connected, the vendor connects to the web interface of the master PLC and begins remote administration of the device and the other subordinate field equipment connected to it. During this time, the master PLC is connected to the internet and has an IP address that can be accessed from anywhere, thus allowing any vendor representative to administer support from any location. Moreover, this configuration can facilitate access to offsite remote Marstad engineers, requiring access to control system operations.

The protocol for establishing this remote access path is manual and unrefined and connections are created on an as-needed basis. The procedure does require that once the remote support activity has been completed the connection is terminated. However, recent increases in demand for water operations has created a need for the configuration to stay enabled, as the cost of terminating and re-establishing the remote access channel becomes quite high and is a union-only function. During the last several months, many cyber incidents were observed and many system abnormalities were attributed to the possibility of cyber attackers exploiting the remote access connection. Management has become concerned that an attack on the Marstad control system could be significant and could impact that ability for the municipality to perform water/wastewater operations in a safe and resilient manner.

The second risk to security is based on the fact that corporate and control system networks are not segmented physically and logically. To accommodate the requirements that meet business logistics and financial operations, the municipal network was developed 'flat' in nature. Although the control system for municipal water operation was originally separated from any other network, the demands associated with the growth in population required that Marstad corporate functions have oversight into infrastructure operations. As such, both the corporate and the control systems operations LAN were connected and now technically reside on the same network. The security is performed by trivial access control listing and some rudimentary routing and ICT operations are the authoritative entity in charge of networks for both corporate and control systems operations.

Although users in both domains are considered trusted, a user from the municipality's city hall or a user from a building outside the control system network can route to areas within the water operations control system network. A main firewall at City Hall provides internet access for all Marstad divisions such as human resources, emergency operations, city planning and transportation. Inside the municipal network, however, the control system network for water and wastewater is not from other networks and no other firewalls provide cyber protection among information enclaves.

Remote access is provided by VNC. Within the main municipal network, VNC is used extensively for remote terminal access. To provide for interoperability, any servers, workstations, or devices that are running VNC on any network are accessible from any IP address within the main city hall network. Passwords on VNC servers are set, but many of the passwords are trivial or shared and an unenforced password policy exists for updating user credentials. The network is essentially flat and the compromise of one corporate machine will make routing around and the compromise of other machines fairly simple.

The two organisations which share water supply with Marstad are connected directly to the control system network to share usage data. Compromise at either location, where security is not managed by Marstad, could potentially create a compromise of Marstad systems. Similarly, a compromise at Marstad could result in a compromise at one of the partner organisations, which could lead to legal liability.

Existing remote access is not managed with effective security policies and procedures. Remote access is not centrally managed and the administrators of the different networks (including the control system network) can create their own remote access solutions. Some users, regardless of role, are authenticated fully to the entire information infrastructure. The data and network security are non-standardised, with connections using very weak encryption or no encryption at all.

Having provided an overview of the Marstad infrastructure, the recommendations can be reviewed using the principles discussed in this practice guide.

Solutions for securing the municipality's remote access

In order to create a robust remote access solution for the municipal water operations, a list of all essential control system devices, processes and services was created. When completed, this list helps asset owners understand what the risk exposure was to control system operations and positions it in terms of the impact an incident could have on core operational services. Using that list, critical assets were assessed to determine what security capabilities existed within them and how they align with the current countermeasures in place to protect those assets.

The risk of exposing the master PLC was considered very high and the procedure was changed to not allow any device connected to the control system to have direct connections to the internet. This, of course, had an impact on the mechanism that facilitated timely vendor support. However, operations had decided that the risk associated with unsecured remote access could render the vendor channel unavailable. In addition, an increased risk for cyber attack was either coming from or going to the vendor and that risk was unacceptable.

For vendor access, a separate authentication terminal server was created and roles and access control lists were created on the server for each and every vendor who needed access. In essence, a dedicated authentication server farm was established to validate credentials and access rights to remote users in the vendor group. Because a problem existed with leaving vendor access channels open, the vendor IDs were automatically disabled after two hours of inactivity or three failed login attempts. The credentials were not removed from the authentication server, however. The new procedure for reactivating a vendor ID was to go through a change control procedure that ordered the reinstatement of the ID and sign-off was required from both sponsoring IT personnel and control system engineers. This new process enhanced the ability to audit vendor connectivity and as each vendor supported different technologies, the Marstad security group was able to create specific intrusion detection signatures that would trigger in the event a deviation occurred from expected traffic flows.

Operational equipment was evaluated in terms of criticality and because of the nature of polling and sampling requirements in the water operations (not high sampling or polling rates by any means), remote access channels were encrypted using a simple commercial off-the-shelf solution. Rather than place the security responsibility of the remote access path on Marstad, this solution was standardised and provided adequate foundation for each vendor to be held accountable for managing their secure remote access. Following this deployment that satisfied the need for external remote access connections to be made, Marstad operations created a terminal server that would facilitate remote access for internal users.

The next step for the municipality to secure their remote connections was to create a clear segmentation between the control systems network and all other networks that were connected to City Hall. A firewall was installed to protect the main router and the multi-homed solutions provided ample opportunity for the segmentation of municipal departments. To compartmentalise the water operations, VLANs were created using dedicated technology. With these segregated networks, only specific VLANs were allowed to access the router that connects to the control system network. No other traffic was allowed to route around the new control system firewall.

Similarly, the two partner organisations were moved off the control system WAN and connected through a separate WAN. This WAN terminated at a new Firewall dedicated for the purpose, which restricts access to all but a usage data stream continuously sending and receiving data between all the partners.

Following the deployment of dedicated firewalls and VLANs, the next step was to create a physical and logical DMZ that was the intermediary between the city hall network and the control systems network. All servers that needed to be accessible to the city hall network were placed into the DMZ. The newly built vendor terminal server was placed in the DMZ as well as the new internal terminal server built solely for internal users. A secure client was then installed on all end devices that could meet the minimal requirements of the software. All VNC software was removed and the secure client software was placed instead on each device that required access from the new internal terminal server. Rule sets for only the necessary ports for the new terminal servers, along with any other servers that were on the DMZ were created (Web servers, SCADA servers, Historians). All other ports and services were removed from the DMZ servers and all users of the core server operating systems were removed (with the exception of the users who manage the server).

The concept of least privilege was used on the DMZ servers and all users managing the servers would have read-only access. Significant improvements were made regarding change control and audit and administrators would have to escalate their privilege with a separate user ID and password in order to make a change on the system.


A full review of all the security policies and procedures was completed and assessed against the control system goals and remote access requirements. The new policy deemed that while using remote access, no other corporate wide services should be required. Access control lists and program limits were implemented to prevent rogue activity. Any remote access to the control systems architecture was considered totally separate and unique from access to only the corporate network. A limited amount of people have network access to the internal terminal server and fewer still have IDs on the terminal server to manage systems on the control network.

A new separate remote access policy was created specifically for the control network. A security requirement was initiated, resulting in the distribution of hardened laptops that would be dedicated for control system remote access. Installed on the laptop were only minimal services, a kernel locking program to prohibit administrative changes in access functionality. The system

was also locked to allow the functioning of only those ports required by the software to manage the control system.

Secure software used to create a VPN full tunnel to the corporate firewall was standardised. In order to use the software remotely, a secure token must be used with one-time passwords along with another password that only the user/operator knows. The encryption on the connection from the software to the corporate authentication server on the corporate DMZ was 256 bits and when the remote access computer was off or in standby mode, the entire disc was encrypted. This solution protects the confidentiality and integrity of the data in transit from an external connection to the corporate network as well as the data at rest. The industrial control systems users were placed on a separate VLAN that only has access to the internal terminal server. The vendor terminal server was also on a separate VLAN. All auditing was turned on at the terminal server and a control system-specific intrusion prevention system was installed outside and inside of the new control system firewall.

Finally, in order to monitor for improper activity, intrusion detection systems were installed behind the network.



(Figure 12: Marstad water distribution network - after security updates)

Conclusion

This document illustrates the methods used to secure remote access into control system networks and demonstrates that existing security techniques deployed in ICT environments are also effective in ICS environments. ICS security staff and ICT security staff each have insights that can be leveraged in either environment. It means that ICS managers do not have to reinvent security procedures and technologies - they are already available.

Modern ICS environments include many of the same technologies as ICT environments and, therefore, share similar security concerns and solutions. Whether an ICT or ICS environment, or a secret facility or a school or a space shuttle, security consists of:

- Preventing unauthorised access to something;
- Preventing unauthorised changes to that something;
- Ensuring that something is available when needed.

The solutions that ensure that these properties of the environment are maintained are the same in both types of environments:

- Direct preventative controls;
- Direct detection and response controls;
- Deterrent controls;
- Compensating controls.

This document has provided many examples of these controls, how they are deployed and managed and to what extent they limit operational risks. A proper regimen of identity management, access controls and active monitoring of networks and systems for malicious activity provide a complete defence-in-depth.

The major difference between security in ICT and security in ICS is the pre-eminence of availability as an operational requirement in ICS environments. ICS environments typically have no tolerance for unplanned downtime, whereas ICT environments, with some notable exceptions, are usually more forgiving. The inability to tolerate downtime in ICS environments results in ICS personnel becoming wary of new technologies; updates to existing systems; changes to the system configurations or the ICS architecture, and so on. However, this is a cultural issue, not a technological hurdle. The requirement for high availability simply requires more planning, more testing and more resources to ensure that security techniques and procedures can be injected into ICS environments safely and without mishap. Sometimes a direct control is not possible, so compensating controls are instead put in place.

This process, wherein one assesses risk, remediates risk, accepts residual risk and continuously improves, is known as a security program and is identical for both ICT and ICS environments. By borrowing from ICT security programs and technologies, ICS managers can appropriately manage their security risks (including remote access risks) without unduly impinging on their operational requirements.

In conclusion, no single secure remote access solution is applicable to all possible architectures and no single remote access solution can provide adequate security without a defence-in-depth approach. However, by exercising caution and generating and implementing concise requirements based on good analysis, secure remote access solutions can be deployed and maintained.

Glossary

Acronyms

AES	Advanced Encryption Standard
ATM	Asynchronous Transfer Mode
CPNI	Centre for the Protection of the National Infrastructure
DHS	U.S. Department of Homeland Security
DMZ	Demilitarized Zone
HTTPS	Hyper Text Transfer Protocol (Secure)
ICT	Information and Communication Technology
ID	Identification
IPSec	Internet Protocol Security
ISDN	Integrated Services Digital Network
IT	Information Technology
L2TP	Layer 2 Tunnelling Protocol
LAN	Local Area Network
MPLS	Multi Protocol Label Switching
NIST	National Institute of Standards and Technology
PC	Personal Computer
PLC	Programmable Logic Controller
POTS	Plain Old Telephone Service
RAS	Remote Access Server
RTU	Remote Telemetry/Terminal Unit
SCADA	Supervisory Control and Data Acquisition
SSH	Secure Shell
SSL	Secure Sockets Layer
VLAN	Virtual Local Area Network
VNC	Virtual Network Computing
VPN	Virtual Private Network
WAN	Wide Area Network
WLAN	Wireless Local Area Network

Nomenclature

Access control	The protection of system resources against unauthorised access; a process by which use of system resources is regulated according to a security policy and is permitted by only authorised entities according to that policy.
Asset	Anything that has value to the organisation, its business operations and its continuity.
Assurance	The attribute of a system that provides grounds for having confidence that the system operates such that the system security policy is enforced.
Authenticate	To verify the identity of a user, user device or other entity, or the integrity of data stored, transmitted or otherwise exposed to unauthorised modification in an information system, or to establish the validity of a transmission.
Authentication	The process of authenticating (see 'Authenticate' above).
Authorisation	The process of determining an entity's rights to perform some action or gain access to a system resource or data set. Alternately, the right or permission that is granted to a system entity to access a system resource (i.e. 'my authorisation is...').
Availability	A property of a system, process, resource or asset that indicates its ability to fulfil its function over a stated period of time, or at a given point in time.
Code of connection (CoCo)	An agreement on the policy and rules for the connection of internal or external assets that are subject to different management domains.
Defence-in-depth	A number of different controls used to mitigate a specific risk (or set of related risks), in order to reduce the residual risk to an acceptable level.
Discretionary access control	A means of restricting access to objects based on the identity and need-to-know of the user, process and/or groups to which they belong. The controls are discretionary in the sense that a subject with certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject.

Encryption	The transformation of plaintext (also called clear- text or that which is in an understandable format) into cipher text (unreadable format) using a method that is difficult to decipher without specific, independent information used in the transformation process.
Impact	The result of an information security incident, caused by a threat, which affects assets.
Least privilege	This principle requires that each subject in a system be granted the most restrictive set of privileges (or lowest clearance) needed for the performance of authorised tasks. The application of this principle limits the damage that can result from accident, error or unauthorised use.
Likelihood	This is broken down into three factors: capability (degree to which the attacker possesses the skills, knowledge and resources to be successful in attempt), opportunity (combination of access to organisation's assets together with vulnerability of environment) and motive/intent (measure of the determination to carry out the attack).
Mitigation	Limitation of the negative consequence of a particular event.
Motivation	A combination of proven intent to attack and the attractiveness of the target in meeting the aspirations and aims of the adversary.
Risk	The potential that a given threat will exploit vulnerabilities of an asset or group of assets and thereby cause harm to the organisation.
Risk assessment	Assessment of threats to, impacts on and vulnerabilities of information and technology processing facilities and the likelihood of their occurrence
Risk management	The process of identifying, controlling and minimising or eliminating risks at an acceptable cost. In the context of this document, risks include security risks that may affect information systems or industrial control systems.
Security level	The combination of a hierarchical classification and a set of non-hierarchical categories that represents the sensitivity of information.
Threat	A potential cause of an incident that may result in harm to a system or organisation.
Vulnerability	A weakness of an asset or group of assets that can be exploited by one or more threats.

Figures

Figure 1: Traditional isolation of corporate and control domains	7
Figure 2: Integrated networks.	8
Figure 3: Attack vectors when using remote access.	24
Figure 4: Example of poor network security.....	28
Figure 5: Security countermeasures.	31
Figure 6: Basic industrial control systems network with no security.....	36
Figure 7: Eliminate 'direct connections' by adding firewalls around the industrial control systems environment.....	38
Figure 8: Add a Control Systems' DMZ and segregate field devices from industrial control systems network.	39
Figure 9: VPN Remote Access.....	41
Figure 10: Secured Remote Access.....	48
Figure 11: Marstad Water Distribution Network - prior to security.	51
Figure 12: Marstad Water Distribution Network - after security updates.....	55

Tables

Table 1: Different remote access solutions.	19
Table 2: Comparison of control systems and ICT remote access requirements.....	33

Further reading

NIST guide to industrial control systems security csrc.nist.gov/publications/PubsDrafts.html#SP-800-82

NIST guide to enterprise telework and remote access security
csrc.nist.gov/publications/nistpubs/800-46-rev1/sp800-46r1.pdf

Mitigations for security vulnerabilities found in control system networks
csrp.inl.gov/Documents/MitigationsForVulnerabilitiesCSNetsISA.pdf,

Common cyber security vulnerabilities observed in DHS industrial control systems assessments
www.uscert.gov/control_systems/pdf/DHS_Common_Vulnerabilities_R1_08-14750_Final_7-1-09.pdf,

Common control system vulnerabilities

csrp.inl.gov/Documents/05-00993%20r0%20Common%20Vulnerability.pdf,

Cryptographic protection of SCADA communications, part 1: background, policies and test plan
www.aga.org/NR/rdonlyres/B797B50B-616B-46A4-9E0F-5DC877563A0F/0/0603AGAREPORT12.PDF

Catalogue of control systems security: recommendations for standards developers

www.uscert.gov/control_systems/pdf/Catalog_of_Control_Systems_Security_Recommendations.pdf

Cyber security procurement language for control systems

www.uscert.gov/control_systems/pdf/SCADA_Procurement_DHS_Final_to_Issue_08-19-08.pdf

Process control and SCADA security

www.cpni.gov.uk/Docs/Overview_of_Process_Control_and_SCADA_Security.pdf

- **Security and SQL Attacks**

Attack methodology analysis: SQL injection attacks (Abstract)
csrp.inl.gov/Documents/SQL%20Abstract.pdf

- **Security and OPC/DCOM**

Understanding OPC and how it is deployed
csrp.inl.gov/Documents/OPC%20Security%20WP1.pdf, Web site last accessed September 2009.

Hardening guidelines for OPC hosts csrp.inl.gov/Documents/OPC%20Security%20WP3.pdf,

Security implications of OPC, OLE, DCOM and RPC in control systems (Abstract)
csrp.inl.gov/Documents/OPC%20Abstract.pdf,

Using operational security (OPSEC) to support a cyber security culture in control systems environments

csrp.inl.gov/Documents/OpSec%20Rec%20Practice.pdf,

Creating cyber forensics plans for control systems csrp.inl.gov/Documents/Forensics_RP.pdf,

Patch management for control systems

csrp.inl.gov/Documents/PatchManagementRecommendedPractice_Final.pdf

- **Modems**

Securing control system modems

csrp.inl.gov/Documents/SecuringModems.pdf,

CPNI good practice guide on firewall deployment for SCADA and process control networks

www.cpni.gov.uk/docs/re-20050223-00157.pdf, *Backdoors and Holes in Network Perimeters: a case study for improving your control system security*

www.us-cert.gov/control_systems/pdf/backdoors_holes0805.pdf

NIST firewall guide and policy recommendations csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf

- **Wireless**

Guide for securing ZigBee wireless networks in process control system environments

csrp.inl.gov/Documents/Securing%20ZigBee%20Wireless%20Networks%20in%20Process%20Control%20System%20Environments.pdf,

Securing wireless VLANs with 802.11

csrp.inl.gov/Documents/Wireless%20802.11i%20Rec%20Practice.pdf,

- **Cyber Security Standards**

A comparison of oil and gas segment cyber security standards

www.us-cert.gov/control_systems/pdf/oil_gas1104.pdf,

A comparison of electrical sector cyber security standards and guidelines

www.us-cert.gov/control_systems/pdf/electrical_comp1004.pdf,

- **NSA Defence-in-Depth**

NSA Defence-in-depth

www.nsa.gov/ia/files/support/defenceindepth.pdf

- **Intruder Detection**

Intruder detection checklist

www.us-cert.gov/reading_room/intruder_det_check.html

- **Personnel Security Guidelines**

Personnel security guidelines

www.us-cert.gov/control_systems/pdf/personnel_guide0904.pdf

Risk assessment for personnel security guidelines
www.cpni.gov.uk/Docs/Risk_Assessment_Pers_Sec_Ed_2.pdf

Personnel security – threats, challenges and measures
www.cpni.gov.uk/Docs/Per_Sec_TCM_v2.pdf

References

1. *Guide to Enterprise Telework and Remote Access Security*, NIST, Special Publication 800-46, Rev.1, csrc.nist.gov/publications/nistpubs/800-46-rev1/sp800-46r1.pdf, Web last accessed December 7, 2009.
2. Williams, TJ, *Purdue Model for Control Hierarchy*, ISBN 1-55617-265-6, 1992.
3. DHS, *Securing Control System Modems* csrc.inl.gov/Documents/SecuringModems.pdf, January 2008, Web last accessed December 10, 2009.