

Table of Contents

Overview	2
Design principles	3
Implementation principles.....	5
Protocol (Procedure).....	6

Overview

External parties may need to connect remotely to critical infrastructure control networks. This access is to allow the manufacturers of equipment used in Australia's critical infrastructure the ability to maintain the equipment, when a fault is experienced that cannot be fixed in the required timeframe any other method. Such access to external parties is to be considered an extraordinary event, and will only be given at critical times where granting access is required to maintain the quality of everyday life in Australia.

Connecting remotely to a computing system is a widely used and well understood task. Cyber-security considerations for such a task can be found in these documents:

<https://www.asd.gov.au/publications/protect/using-remote-desktop-clients.htm>

https://www.asd.gov.au/publications/protect/additional_security_considerations_and_controls_for_vpn.htm

Connecting remotely to a control system has some specific considerations. There is existing literature on the topic of remote access to control systems, such as international standard IEC 62443 and advice from ICS-CERT:

<https://ics-cert.us->

[cert.gov/sites/default/files/recommended_practices/RP_Managing_Remote_Access_S508NC.pdf](https://ics-cert.us/sites/default/files/recommended_practices/RP_Managing_Remote_Access_S508NC.pdf)

This document is broken into three sections:

1. **Design principles.** The design principles include topics such as time limiting the connection, strong authentication, and the creation of well managed devices.
2. **Implementation principles.** The implementation principles provide guidance on good approaches for satisfying the design principles.
3. **The protocol.** Once the design and implementation principles have been followed, the specified protocol, or procedure, for remote access may be followed.

This document should be ratified every six months to ensure the incorporation of any necessary updates due to a changing cyber-threat landscape. Further, if a significant cyber-event, or awareness of a new method or tradecraft, happens outside of this review cycle, adjustments to the process outlined in this document should occur and be publicised immediately.

Design principles

1. By default, there should be no communication between the vendor and the critical infrastructure control system.
2. Networks should be segmented and segregated. Details on the application of segmentation and segregation can be found here:
https://www.asd.gov.au/publications/protect/network_segmentation_segregation.htm
Included in this process will be firewalls. The firewalls should be configured as tightly as possible, including restricting to specific protocols, ports, MAC and IP addresses, and directions. For example, a stateful firewall could be employed which allows communication to be initiated only in one direction. Internet facing firewalls and the control system DMZ firewall should be completely separate devices.
3. There must be other processes and procedures in place before this protocol is used. These processes include:
 - a. a way to disconnect the control system from the internet quickly, if unwanted external control or actions are detected;
 - b. a way to revert the control system to a known good state, if an unwanted or unauthorised configuration change is detected;
 - c. a cyber-incident response plan in case malware is introduced; and
 - d. the expected safety plans, in case an unwanted physical action took place.
4. Multi-factor authentication should be used. Two-factor authentication should be used at a minimum. Details, including why this is needed, can be found here:
https://www.asd.gov.au/publications/protect/multi_factor_authentication.htm.
5. Ensure the login credentials are such that a specific person at the remote end is attributed to the actions, rather than a generic login for an organisation. Ensure these person-specific details are recorded, for example in an Access Log or Engineering Change Request.
6. Time limit the connection, e.g. to 24 hours or the length of a shift, and ensure the credentials are one-time-use credentials. Ensure the credentials expire after 24 hours whether they are used or not.
7. If the connection is inactive for more than 30 minutes, the connection should be removed. 30 minutes may be restrictive, if processes essential to the task such as firmware updates take longer than 30 minutes. If longer is required, a case should be put forward and relevant records kept.
8. Ensure there is a procedure to acquire approval for connection of remote access by a senior officer of the organisation. If there are particular necessary notifications in various jurisdictions, list them here.
9. Ensure the device used at the remote (vendor) end is used solely for the purpose of connecting to the Australian critical infrastructure organisation. That is, the computer at the remote organisation cannot be used to connect to country X yesterday, country Y the day before, Australia today, and country Z tomorrow. The preference is to use a laptop at the vendor's end provided by the Australian Critical infrastructure organisation, rather than relying on third party infrastructure. The computer should only be used for connecting to Australia's critical infrastructure organisations, for the purpose of accessing the control system once through the various internal connections. One computer at the remote organisation dedicated to one Australian organisation is the ideal, and if this ideal cannot be met explain how you are going to mitigate the risk.

10. Apply ASD's "Top 4" to the highest maturity level, with the rest of the "Essential 8" where applicable (detailed here: https://asd.gov.au/publications/protect/top_4_mitigations.htm and <https://asd.gov.au/publications/protect/essential-eight-explained.htm>) on the computer at the remote end. It is noted that the process of managing a standalone computer by the remote organisation will be challenging.
11. Apply ASD's "Essential 8" where applicable (detailed here: <https://asd.gov.au/publications/protect/essential-eight-explained.htm>) on all interim machines internal to the critical infrastructure organisation's network that are not prevented from such measures for OT reasons. Example computers that could have their security improved in this way include machines in de-militarised zones (DMZs) and jump boxes in general, and any machine used by the CI organisation to view, control or supply credentials to the vendor's connection. Examples of what "where applicable" covers is that if Microsoft Office is not installed, then macros may not be an issue; and if no data is locally stored, daily backups may not be unnecessary.
12. "Bastion hosts" (special-purpose computers on a network specifically designed and configured to withstand attacks) and interim machines should be turned off whenever possible, when not in use to prevent attackers acquiring a foothold while the remote vendor protocol is not in use, and waiting for a connection to the internet. Having machines turned off when not in use will add complexity to the requirement to maintain the "Top 4" and "Essential 8", so a plan needs to be created to manage this.
13. To aid in mitigating the risk of supply chain attacks, critical infrastructure operators and vendors should put in place robust mechanisms to verify all software and tools used in the remote vendor access protocol process. This includes the engineering analysis and fault investigation tools.
14. Ensure contractually that any data viewed or acquired as part of the remote access is used only for the purpose of resolving the issue the remote access was granted for, and must be returned to the critical infrastructure organisation and destroyed at the remote access end either when the issue is resolved, or after the period of 1 year, whichever is sooner.
15. Ensure contractually that there is an ability to audit the organisation at the remote end to ensure each of the conditions is met. These conditions include "device only used for Australia", device has had ASD's Essential 8 applied where applicable, data is kept only for as long as it needs to be and a copy of what was obtained is returned to the critical infrastructure organisation, etc. Also explicitly looked for should be the applicable cipher-suites available at the vendor's end, such that the risk of a down-grade of cipher-suite attack is managed.
16. Ensure contractually that there is ability to periodically red-team test the protocol at all parts of the protocol, including the remote vendor's end.
17. Ensure contractually that any connectivity and hosting requirements for the remote access infrastructure is specified. That is, consider the ability to perform denial of service attacks against the remote vendor protocol, and the impact of this to the critical infrastructure organisation in an emergency.

Implementation principles

1. There should be no connection directly from the vendor to the control equipment. There should be at least one hop through a device controlled and managed by the critical infrastructure owner first, and ideally the ingress and egress from that point should be based on different communication protocols.
2. The connection between the control network and any external device should be physically disconnected when the protocol is not being used. Ideally this would mean removal of a physical cable, however sometimes the location to disconnect the cable may be a significant distance from the control equipment. An option may be a connection en-/disable via key operated switch, located closer to or in the control room.
3. Ensure there is a warning on the control system operator's machines when the connection to the remote vendor is established.
4. VPNs should be configured as tightly as possible, for example locked to specific IP addresses through firewalls.
5. Ensure accounts are unprivileged accounts for as many of the systems as possible. For example, only have administration privileges on the final device in the communication path, if administration privileges are required. Detailed information on restricting administrative privileges can be found here:
https://www.asd.gov.au/publications/protect/restricting_admin_privileges.htm
6. Ensure the connection from the Internet is to a jump box, and not directly to the control network. Information about jump boxes is available in this file:
https://www.asd.gov.au/publications/protect/Secure_Administration.pdf
7. Internal passwords should never be given to the entity performing the remote login. For example, one solution may be that since the critical infrastructure organisation will have an employee attending and witnessing actions, the critical infrastructure organisation's employee could type in the login details whenever passwords are required. Note that some solutions for viewing remote sessions create connections via a server, potentially running at a third party's organisation and hence routing all traffic via the third party organisation, which is not acceptable. Another solution might be to implement user-privilege control.
8. Keep a record of who logged in, when they logged in, what activities are to be performed while logged in, and for what purpose. Keep this record for at least 5 years, unless exceeded by other requirements.
9. Ensure that the connection from the control network to the Internet is well controlled; i.e. only allowed to specific hosts and ports.
10. Ensure there is an inline tap or other capability able to record all traffic through the connection from the Internet. Ensure this recording is at such a place that the traffic is recorded in a format (e.g. plaintext or encrypted with a known key that is also kept as part of the record) so that the traffic can be analysed. Keep this recording securely stored for a minimum of 5 years.
11. Split tunnelling, the ability to connect to different security domains using the one interface such as a network interface card (NIC), should be disabled.

Protocol (Procedure)

1. Ensure procedures discussed in Design Principle 2, such as a method to disconnect the control system from the internet quickly, are known and in place.
2. Obtain approval for the connection from the relevant senior officer in the critical infrastructure company.
3. Enter in a log the data and time (change record), who will be connecting, who the witness will be, and the reason for the connection.
4. Issue the time-limited password (from Design Principle 6) to the remote entity using existing channels.
5. An authorised person employed by the critical infrastructure organisation should attend and witness all actions from the time the cable to the internet is connected, through until the cable is removed.
6. Physically create a path from the control network (be that by connecting a cable, or turning a keyed physical switch) to the network outside of the control network, so as to create a pathway (via the DMZs, jump boxes etc.) to the Internet.
7. Ensure the inline data capture, at a location which can see the traffic in plain text, is operational.
8. Invite remote party to connect.
9. Authenticate the person and the hardware at the remote end. Multi-factor, at least two-factor, authentication should be used to authenticate the person.
10. The remote end should authenticate the critical infrastructure organisation also, i.e. “mutual authentication”.
11. While the external party is connected, an authorised person employed by the critical infrastructure organisation should type in any internal passwords for the connection as required. No internal passwords to be given to the remote entity.
12. Once the work is completed, physically remove the connection to the Internet via either removal of the cable, or turning of a keyed switch. It is expected that the time that the connection to the Internet is in place is measured in hours, not days.
13. Cease the inline data capture and securely store the data captured. The data should be stored for a period of at least 5 years. Best practices for data storage, such as off-site backups, should be observed.
14. Enter into the log the date and time of the completion of the remote access work, and the name and location of the corresponding inline data capture file. Ensure the log is kept for 5 years.