Managing cyber security risks

Report 3: 2019–20

Queensland Audit Office

Better public services

Queensland Audit Office
Better public services

Your ref:
Our ref: PRJ01205

1 October 2019

The Honourable C Pitt MP
Speaker of the Legislative Assembly
Parliament House
BRISBANE QLD 4000

Dear Speaker

Report to parliament

This report is prepared under Part 3 Division 3 of the *Auditor-General Act 2009*, and is titled *Managing cyber security risks* (Report 3: 2019–20).

In accordance with s.67 of the Act, would you please arrange for the report to be tabled in the Legislative Assembly.

Yours sincerely

Brendan Worrall
Auditor-General

Contents

Audit objective and scope

This audit examined whether entities effectively manage their cyber security risks.

We addressed this by assessing whether entities:

- understand and assess the extent to which their information assets and organisational processes are exposed to cyber security risks
- design and implement effective information controls to mitigate identified cyber security risks.

We selected three entities for this audit. We do not want to compromise the security of these three entities by publicly identifying their security vulnerabilities, so we have not named them in this report. We acknowledge the three entities have different levels of resourcing and capability for managing cyber security risks.

We use the term ‘entities’ in this report to refer broadly to all Queensland public sector entities (departments and statutory bodies) and local governments.

Audit approach

Our audit included detailed technical testing by specialist security consultants including:

- a ‘red team’ assessment for each of the three entities. A red team engagement tries to find the quickest method to access an entity’s security mechanisms and compromise its sensitive applications and data. In doing so, it considers the target and resources available, and may attempt social engineering, physical entry, and data exploitation
- an ‘open source threat intelligence assessment’ to determine whether any sensitive information about the three entities could be obtained from publicly available sources
- testing whether the entities have implemented four of the 'Essential Eight' mitigation strategies published by the Australian Cyber Security Centre (ACSC) to help organisations protect their systems against cyber threats.

Scope exclusions

We did not, as part of this audit, examine the effectiveness of activities conducted by the Queensland Government Chief Information Office (QGCIO). Only one of the three entities in this audit uses the services of the QGCIO.

Further details about the scope and approach are in Appendix B.

Technical language

Cyber security is a complex, technical field, and the language about it reflects this. It is necessary to use some of this language to be precise, but we have explained and simplified it in this report and provided definitions when necessary. When a term is very complex, we have included more detail in the glossary.

Reference to comments

In accordance with s. 64 of the *Auditor-General Act 2009*, we provided a copy of this report to relevant entities. In reaching our conclusions, we considered their views and represented them to the extent we deemed relevant and warranted.

Introduction

Protecting important information assets with secure systems is critical to Queensland's economic and security interests. The Global Risks Reports produced by the World Economic Forum in 2018 and 2019 found that 'data fraud or threat' and 'cyber attacks' are in the top five most likely global risks in terms of likelihood (along with environmental risks).

A Microsoft-commissioned study by Frost and Sullivan (a research and consulting firm) estimated the potential direct economic loss of cyber security incidents on Australian business as $29 billion per year. When factoring in other indirect costs—such as damage to business reputation and loss of customer base—the actual loss is even higher.

Media reports show an alarming trend of growing cyber security attacks and corporate espionage by foreign state-sponsored hackers and criminals targeting Australian Government entities. These are organised, targeted, deceptive cyber attacks intended to compromise Australia's economic interest, and national security.

The *2017–18 Cyber Security Survey*, conducted by BDO Australia and AusCERT, stated that organisations seeking to enhance their cyber security capabilities will need to get a better understanding of the cyber threats related to them and their industry. The survey report identified the following threat actors (those conducting malicious activities against entities):

- hacktivists—who target computer networks to advance their political or social causes
- criminals—including individuals and sophisticated criminal groups who steal personal information and extort victims for financial gain
- insiders—who typically steal their organisations' information for personal, financial, or ideological reasons
- nation-states—who target systems to steal sensitive state secrets for economic and political advantage.

Better practice frameworks

We used two better practice frameworks to assess the three in-scope entities for this audit—the Queensland Government's *Information security policy (IS18:2018)* (the information security policy) and the Australian Cyber Security Centre's (ACSC) 'Essential Eight' mitigation strategies to help organisations protect their systems against cyber threats.

We acknowledge that it is not mandatory for all entities to apply these frameworks, but these frameworks provide good guidance to all entities on how to effectively manage cyber security risks.

Information security policy (IS18:2018)

The information security policy is aligned with other better practice frameworks for managing cyber security risks—the International Standard for Information Security Management Systems (ISO 27001), the ACSC's Essential Eight strategies, and other supporting frameworks developed by the Queensland Government Chief Information Office.

The information security policy helps entities apply a consistent, risk-based approach to information security in order to safeguard the confidentiality, integrity, and availability of the data and information they maintain.

The policy applies to core Queensland Government departments (as defined by the *Public Service Act 2008*). The Queensland Government Chief Information Office states it has no mandate to require local government bodies to comply with the information security policy. But it strongly encourages entities (including Queensland government owned corporations, universities, and local governments) to do so, to demonstrate better practice.

Appendix C shows the policy requirements of the information security policy.

Mitigation strategies for cyber security risks

The ACSC published the 'Essential Eight' mitigation strategies in 2017 to help organisations protect their systems against cyber threats. On 1 October 2018, in policy requirement three of the information security policy, the Queensland Government Chief Information Office made the Essential Eight mitigation strategies a minimum security requirement. For this audit, we focused on what the ACSC calls the 'Top 4' strategies, because it has stated that, if organisations effectively implemented these, they would mitigate at least 85 per cent of cyber intrusions.

The Top 4 mitigation strategies include:

- **application whitelisting**—controls to block all non-approved or malicious applications from being executed in an information and communication technology (ICT) environment
- **patching applications**—controls to address known vulnerabilities in the security of applications that can be exploited by threat actors executing malicious code
- **restricting administrative privileges**—controls to minimise the risk of threat actors exploiting privileged system access (which is held by users who can access sensitive data and create and configure within the system)
- **patching operating systems**—controls to address known vulnerabilities in the security of operating systems that can be exploited by threat actors executing malicious code.

Summary of audit findings

Identifying and assessing cyber security risks

Frameworks for managing cyber risks

Two of the three entities we audited had established governance arrangements for managing cyber security. Their actions demonstrated management's commitment to an effective information security culture in their business. The third entity was starting to place more emphasis on managing cyber security risks but had not yet established an effective framework for doing this.

We found the following elements are essential for having an effective cyber security framework:

- defining a high-level approach for managing information security
- having an information security policy that defines an entity's objectives for managing information security
- incorporating information security within an entity's corporate governance
- implementing mandatory information security training for all staff
- establishing an information security management team
- requiring periodic reporting from the entity's chief information officer on the current security threat level, identified vulnerabilities, and progress on actions to mitigate cyber security risks
- disciplining users who commit breaches in information security
- defining the entity's appetite (the risk an organisation is willing to take in order to meet its objectives) for information security risk at the enterprise and operational level
- allocating specific responsibility for coordinating an entity's risk assessment process for cyber security across all organisational units.

Managing key information assets

The Queensland Government Information Security Office updated its *Queensland Government Information Security Classification Framework* (QGISCF) in 2018. The QGISCF recommends that entities classify their information assets according to business impact and implement appropriate controls according to the classification. (This is mandatory for core Queensland Government departments.)

None of the three entities had effectively implemented a process for applying a security classification to information assets, but all three had plans to address this.

They had not yet conducted a comprehensive assessment to ensure they had identified all their information assets that were at risk. Nor had they considered what controls they had or what was necessary to protect those assets.

We also found the three entities did not have a full record of the ICT assets they allocate to their employees. Their processes for managing employee separations (for example, resignations, retirements, and dismissals) were not robust enough to ensure the entities knew all employees returned their ICT assets.

For two of the entities, we found almost 750 ICT assets (according to their records) were assigned to employees who no longer work for them. Either the entities' asset records are out of date, or there is a risk that these assets could be used to access the entities' sensitive information.

Identifying and assessing cyber security risks

One of the three entities periodically assesses its exposure to cyber security risks at an enterprise level and for business-critical applications. One of the entities has started to do this, but the other did not have a risk assessment process to identify and assess cyber security risks.

While the third entity registered some general cyber security risks in its risk registers, it described them at a corporate level only. This means the entity's risk treatments may not be enough to address risks to specific information assets.

Mitigating cyber security risks

Application whitelisting

None of the three entities had a strategy for implementing application whitelisting, but one has started to implement it in its server and desktop environment. We observed two challenges entities have in implementing application whitelisting:

- One of the entities was concerned about making sure it provided its users with a collaborative working environment, while at the same time protecting sensitive information. Entities in this position should at least consider implementing application whitelisting on devices used to access sensitive information.
- Implementing application whitelisting is complex because of the multiple layers of technology to which it needs to be applied
  - One of the entities implemented application whitelisting on its Windows 10 desktops and newer servers. It didn't implement it on its servers using older technology (which accounted for 93 per cent of servers hosting applications) or its desktops on older versions of the Windows operating system (which accounted for 33 per cent of its desktop computers).
  - It also allowed Windows 10 users to run scripts (sequences of instructions interpreted within a program) based on their access privileges, which means privileged users could run scripts from an untrusted source. A potential attacker could exploit this weakness if they could compromise a privileged user account.

Patching operating systems and applications

One of the three entities demonstrated effective controls for patch management, including:

- having a strategy for patch management
- having effective processes for implementing patches
- effectively prioritising the implementation of patches assessed as extreme risk, and adequate processes for implementing patches assessed as below extreme risk
- replacing/updating legacy systems (older systems for which there is no longer any support from the supplier/vendor) to vendor-supported versions
- mitigating vulnerability risks through its risk and governance processes when patches are not available.

One of the other entities had developed patch management procedures, which now need to be implemented throughout its organisation, while the other entity's approach for patching systems was ad hoc and not defined in a documented process.

Restricting administrative privileges

We found all three entities had implemented effective controls to minimise the number of people with database administrative privileges. But only one of the three entities had implemented effective controls for managing privileged user access for the operating systems. It had processes for assigning, reviewing, and making privileged users re-apply for access after 12 months.

For the other two entities:

- One had designed and implemented a process for administering and managing privileged use access; however, some business units that operate their ICT environments independently did not apply a robust process for managing administrator privileges.
- One had not documented its process for administering and periodically reviewing privileged user access. We found evidence that a privileged user no longer required that access for their current role.

None of the three entities had fully implemented controls to ensure administration tasks could only be performed through a secure connection. This creates a risk that if an attacker gained access to their networks and compromised a privileged user account (administrator account), they wouldn't need any further authentication beyond the initial compromised password.

Privileged user accounts should not be used for business-as-usual activities such as reading email and web browsing. This is to reduce the risk of a privileged account being used to download malicious software or being subject to a phishing attack. ('Phishing' refers to fraudulent scamming attempts to obtain sensitive information from users.)

In terms of implementing controls to prevent privileged users from reading emails, opening attachments, browsing the web, or obtaining files from internet services like instant messaging and social media:

- One entity had implemented effective controls.
- One had not implemented controls to prevent users from downloading malicious content.
- One had implemented some controls around email, but still provided privileged users with unrestricted access to the internet (which includes webmail).

In terms of logging and monitoring privileged user accounts:

- One entity had implemented effective controls to log and monitor privileged user activity.
- One kept logs of privileged user account activities but did not monitor or receive alerts of privileged user account activities.
- One had implemented limited logs of user activities in general, but this was not targeted to privileged users. The entity did not monitor the logs or receive alerts on anomalies.

Managing security risks in the supply chain

All three entities need to improve their practices for managing security risks with their suppliers in their supply chain. For example:

- One entity did not define the information security-related roles and responsibilities within its general procurement process for managing information security risks in its supply chain.
- One entity did not regularly review and monitor third-party ICT supplier services to ensure they maintain appropriate security levels. It did not have a risk assessment process to determine the suitability of potential external service providers from an information security perspective.
- One entity used standard contract clauses in its contracts with third-party providers, which do not include specific provisions for information security with which the vendor must comply. There was no process for monitoring, reviewing, or auditing the vendor's compliance with information security requirements.

Cyber security awareness training

Two of the three in-scope entities provided cyber security awareness training to all staff, and one of these entities provided more targeted training to users who have access to sensitive data.

We found the third entity did not provide any cyber security awareness program. This entity is therefore more susceptible to the type of attacks that take advantage of people being a weak link in the chain of defence.

Security testing

For each of the three entities, we nominated a narrow set of targets for our security consultants to test, based on our understanding of the entities' key information assets.

In all three instances, our security consultants were successful in compromising at least one target we set for each entity. We have advised each of the in-scope entities of the risks we identified (specific to each entity) through our testing. We acknowledge their efforts since our audit and their ongoing plans to mitigate the risk of cyber security attacks against their sensitive information.

Our security consultants found the following issues can make it easier to compromise a target, and we provide these as learnings for all entities:

- physical security—challenging someone who tries to access an office facility without authorisation is important. It prevents an attacker from obtaining direct access to the entity's internal assets and increasing their ability to attack the entity's sensitive data
- password practices—easily guessable passwords make it simpler to compromise user accounts and use these accounts to gain control over an entities' networks. Common passwords such as ‘welcome’, ‘password’, and ‘newuser’ make it easier for an attacker to gain access to an entity’s systems
- known password breaches—if users use their corporate email address on online services and have the same password as they do on their corporate network, an attacker could use breached user accounts and passwords (which are publicly available on several sites) to access an entity’s network’
- multi-factor authentication—should be used to prevent users from remotely logging into an entity’s internal network without requiring two-factor authentication (for example, a username and password, plus a code sent to a mobile phone). This should also be used to ensure staff and administrators can only access sensitive internal servers (from within their networks) with multi-factor authentication
- administrative accounts—when users with administrative privileges use the accounts for business-as-usual activities (such as accessing email), attackers find it easier to gain control of an entities' systems once they compromise administrator accounts
- network segmentation—a lack of network segmentation allows an attacker to move laterally within an entity’s networks once they access the internal networks
- outdated systems—entities using systems that are running outdated applications and operating systems that have not been supported by vendors in several years, have an increased risk exposure, because security vulnerabilities are likely to exist and unlikely to be fixed by vendors
- descriptive subdomains—may provide attackers with an insight into an entity’s ICT environment. This could indicate what online services an entity uses or identify non-production environments (for example, through a subdomain name like ‘development.entityname.qld.gov.au’) that may not be as strongly secured as production environments
- insecure encryption channels—when an entity uses online application hosts (a website that allows users to use a software application over the web, for example, a mapping service) without encryption, an attacker could use this to manipulate communication between users and online services.

Audit conclusions

The three entities we audited are not managing their cyber security risks as effectively as they could. One of the entities demonstrated a higher level of maturity in cyber risk management across its governance and technical mitigating strategies than the others. But this was not enough to prevent our security consultants from compromising its ICT environment. The fact that our consultants successfully compromised all three entities' ICT environments and could access their sensitive or non-public data demonstrates there were gaps in their mitigation strategies.

Two of the three entities had appropriate frameworks for managing cyber risks. While they had some elements of effective cyber security processes in place, our testing demonstrated vulnerabilities that need to be addressed. Addressing these vulnerabilities is a balancing act between risk appetite and cost.

None of the three entities could demonstrate an understanding of the extent to which its information assets were exposed to cyber security risks. All three entities need to conduct a comprehensive assessment of their information assets to determine which assets are at risk and require further controls to protect. Without this, it is difficult to know whether an entity has implemented the right level of controls to protect its assets. We recognise that entities must make decisions regarding the extent of controls they will invest in, and that they will never be fully effective in mitigating all risks in an ever-evolving threat landscape.

None of the three entities has effectively implemented the Top 4 mitigation strategies for cyber security risks. This demonstrates that some other entities may also find it challenging to implement this better practice guidance.

As entities use more cloud-based services that provide remote access into their systems, they need to be vigilant in assessing how vulnerabilities in their service providers could expose them to cyber risks.

They also need to make sure their users are aware of their responsibilities in managing cyber risks. In particular, we found poor password practices unnecessarily exposed the three entities to attack. Third-party providers and internal staff could be the weak links in an entity’s line of defence.

Recommendations

We provided each of the three in-scope entities with detailed recommendations relating to the issues we identified relevant to them.

For the benefit of all entities, we provide the following recommendations, drawn from the learnings of this audit.

We recognise that implementing effective controls for cyber security should be performed on a cost-benefit basis. Therefore, we recommend all entities firstly assess themselves against recommendations 1–3, which will help them ensure they have a framework for managing cyber security risks, know what information assets they have, and know to what extent those information assets are exposed to cyber security risks. Then, based on the results of these activities, entities should consider how relevant recommendations 4 to 17 are for their risk appetite and exposure.

All entities

We recommend that all entities self-assess against the findings of this report, and where relevant:

Cyber security framework

1. develop a framework for managing cyber security risks consistent with the *Information security policy (IS18:2018)* (Chapter 2)

   They should also have information security standards to ensure the framework is consistently applied throughout the entity at an operational level.

   Queensland Audit Office (QAO) insight statement 1 in Chapter 2 provides more guidance on this.

Information classification

2. develop and implement policies and procedures to identify and classify information assets, so they can effectively manage all their information assets that are at risk. This should include policies and procedures for:
   - identifying and maintaining an inventory of information assets
   - classifying information assets as per the 2018 *Queensland Government Information Security Classification Framework* (Chapter 2)

Identifying and assessing cyber security risks

3. develop and implement a methodology for identifying and assessing cyber security risks to their information assets. This should include:
   - developing a risk assessment process for cyber security that integrates with their enterprise risk management framework
   - developing risk appetite statements for cyber security
   - identifying and assessing cyber security risks to their key information assets (Chapter 2)

   QAO insight statement 3 in Chapter 2 provides more guidance on this.

Information asset management

4. review how they manage their ICT assets by:
    - reviewing their list of ICT assets and checking if they are assigned to employees who no longer work there and, if necessary, recovering any ICT assets that have not been returned
    - reviewing their employee separation process to ensure it includes updating the ICT asset register whenever an employee's employment ends (Chapter 2)

    QAO insight statement 2 in Chapter 2 provides more guidance on this.
5. assess the adequacy of their physical security to protect their ICT assets from unauthorised access (Chapter 4)

Cyber security risk mitigation strategies

6. design and implement an application whitelisting strategy (Chapter 3)

    QAO insight statement 4 in Chapter 3 provides more guidance on this.
7. design and implement a patch management strategy to cover the patching of vulnerabilities in operating systems, applications, drivers, and hardware devices (Chapter 3)

    QAO insight statement 5 in Chapter 3 provides more guidance on this.
8. ensure they effectively minimise and restrict administrative privileges (Chapter 3)

    QAO insight statement 6 in Chapter 3 provides more guidance on this.
9. implement risk management practices for their use of third parties to deliver information technology services (Chapter 3)

    QAO insight statement 7 in Chapter 3 provides more guidance on this.
10. undertake a risk assessment to determine the most effective password policy and implement it as a priority (Chapter 4)

    Controls may include:
    - blacklisting commonly breached passwords, dictionary words, and words about the context of the work environment (for example, entity name, services, and units)
    - preventing the use of repetitive and sequential characters.

    Better practice guidance that may help entities includes:
    - National Institute of Standards and Technology (NIST) Special Publication 800-63B *Digital Identity Guidelines*
    - *Australian Cyber Security Centre Information Security Manual*
    - Queensland Government Enterprise Architecture Guideline: *Reducing password frustration for Queensland public servants*
11. implement multi-factor authentication as a minimum on external services that allow login with their domain accounts, and for sensitive internal systems (Chapter 4)
12. review all subdomains and consider whether they provide an indication of the entity's underlying technology or services, and modify existing subdomains to obscure exposing information (Chapter 4)
13. implement encryption on online services that communicate via an unencrypted channel (Chapter 4)
14. segregate workstations located in publicly accessible areas from their corporate network (Chapter 4)

15. develop cyber security training and deliver it to all staff, with more targeted training to users who have access to sensitive data (Chapter 3)

    QAO insight statement 8 in Chapter 3 provides more guidance on this.

16. ensure security and awareness training includes:
    - discouraging the use of corporate email addresses on external services
    - education on the risks of posting information on social media that provides information on an entities' technology services
    - education on phishing attacks
    - education on the risk of physically 'tailgating' people into public sector buildings and offices (Chapter 4)

Monitoring and logging

17. introduce and configure end user device logging.

    This should include configuring security logs and rules on end user devices (for example, computer desktops and laptops) for detecting malicious and anomalous behaviour and events. (Chapter 4)

1. Context

This chapter provides the background to the audit and the context needed to understand the audit findings and conclusions.

Information asset security classifications

It is important that entities identify what information assets they own that may be a target of attackers (which may be personal information or information the entities create or obtain in running their operations). When entities do this, they can better target their risk mitigation strategies to their high-risk information assets.

In September 2018, the Queensland Government Chief Information Office published a new version of its information security classification framework. This resulted in changes to the business impact levels for confidentiality to align with the Australian Government approach for classifying information.

The framework explains that, to classify their information assets, entities need to:

- determine the business impact levels of the loss, compromise, and misuse of their information in terms of the impact on confidentiality, integrity, and availability
- analyse their information and information assets against the business impact levels they have created and assign confidentiality, integrity, and availability values
- determine and apply appropriate controls to safeguard the information and information assets in a consistent manner
- regularly assess whether the controls assigned for confidentiality, integrity, and availability values are adequate to keep the entity within its risk tolerance level (the risk it's willing to take in order to meet its objectives).

The business impact levels (designed for state government departments) for confidentiality are:

- **OFFICIAL**—low or negligible confidentiality impact. This classification represents most Queensland Government information by volume but the lowest business impact per document if compromised or lost. All routine public sector business, operations, and services are treated as OFFICIAL.
- **SENSITIVE**—moderate confidentiality impact. This information requires additional care due to its sensitivity or moderate business impact if compromised or lost. Examples include personal information, legal professional privilege, and government or agency business whose compromise could affect (1) the government's capacity to make decisions or operate, (2) the public's confidence in government, or (3) the stability of the marketplace.
- **PROTECTED**—high confidentiality impact. This information requires the most careful safeguards due to its sensitivity or major business impact if compromised or lost. PROTECTED information assets require a substantial degree of control as compromise could cause serious damage to the state, the government, commercial entities, or members of the public.

When an entity has determined high confidentiality information to be at the PROTECTED level, it must consider the PROTECTED controls outlined in the current information security manual published by the Australian Cyber Security Centre (ACSC).

Mitigation strategies for cyber security risks

The ACSC has compiled a list of mitigating strategies entities can use to improve their ability to protect against cyber security risks. It has developed eight mitigation strategies it says should be implemented as a baseline where practicable.

The 'Essential Eight' strategies are explained in Figure 1A.

Figure 1A
'The Essential Eight'

Mitigation strategies	Purpose
Application whitelisting of approved/trusted programs to prevent execution of unapproved/malicious programs.	All non-approved applications (including malicious code) are prevented from being run and installed.
Patch applications Patch/mitigate computers with 'extreme risk' vulnerabilities within 48 hours. Use the latest version of applications.	Security vulnerabilities in applications can be used to execute malicious code on systems.
Configure Microsoft Office macro settings to block macros (a series of commands and instructions grouped together to automate a task) from the internet, and only allow vetted macros either in 'trusted locations' with limited write access or those that are digitally signed with a trusted certificate.	Microsoft Office macros can be used to deliver and execute malicious code on systems.
User application hardening Configure web browsers to block Flash (ideally uninstall it), advertisements, and Java on the internet. Disable unneeded features in Microsoft Office, web browsers, and PDF viewers.	Flash, ads, and Java are popular ways to deliver and execute malicious code on systems.
Restrict administrative privileges to operating systems and applications based on user duties. Regularly revalidate the need for privileges. Don't allow privileged accounts to be used for reading email and web browsing.	Administrative accounts are the 'keys to the kingdom'. Adversaries use these accounts to gain full access to information and systems.
Patch operating systems Patch/mitigate computers (including network devices) with 'extreme risk' vulnerabilities within 48 hours. Use the latest operating system version. Don't use unsupported versions.	Security vulnerabilities in operating systems can be used to further compromise systems.
Multi-factor authentication including for remote access and for all users when they perform a privileged action or access an important data repository (one that contains sensitive information or information that needs to be highly available).	Stronger user authentication makes it harder for adversaries to access sensitive information and systems.
Daily backups of important new/changed data, software and configuration settings, which is then stored, disconnected, and retained for at least three months. Test restoration of backups initially, annually, and when ICT infrastructure changes.	To ensure information can be accessed again following a cyber security incident (for example, a ransomware incident).

Source: Australian Cyber Security Centre, Australian Signals Directorate, cyber.gov.au/publications/essential-eight-explained.

For this audit, we focused on what the ACSC calls the ‘Top 4’ strategies, because it states that if organisations effectively implemented these, they would mitigate at least 85 per cent of cyber intrusions. The Top 4 strategies include application whitelisting, patching applications and operating systems, and restricting administrative privileges.

2. Identifying and assessing cyber security risks

This chapter is about how effectively the three entities we audited identify and assess their exposure to cyber security risks.

Introduction

An entity needs to understand and assess its exposure to cyber security risks. It can then develop appropriate mitigation strategies to reduce the risk of its information assets being compromised through a cyber security incident.

To assess effectiveness, we examined whether the three in-scope entities for this audit:

- design adequate frameworks to identify and manage their cyber security risks
- clearly identify and manage their key information assets
- regularly perform risk assessments of their information security to effectively identify, manage, and understand their cyber security risks.

We have created a series of insight statements to summarise the effective practices we noted during the audit. Our intention is that other entities can use these statements to assess the adequacy of their practices.

Frameworks for managing cyber security risks

Two out of the three in-scope entities had established information security governance arrangements and frameworks for managing cyber security risks. Their actions demonstrate their management's commitment to an effective information security culture in their business environment.

We found the third entity was starting to place more emphasis on managing cyber security risks. It had not, however, established an overarching information security policy and framework to set the objectives for its information security and its approach for managing these objectives.

The third entity had no specialist to manage information security on a day-to-day basis. It relied on members of its information technology group to be responsible for information security in addition to their operational responsibilities—but their roles in relation to information security were not defined. They were limited by their existing operational commitments, which often took priority over their information security responsibilities. This means the entity was prone to being reactive to security incidents rather than planning for mitigation.

QAO insight statement 1 shows important elements in an effective framework for managing cyber security risks.

QAO insight statement 1

Establishing a framework to manage cyber security risks

Entities need to establish a framework and governance arrangements for cyber security to ensure they have the right control environment and culture regarding the risks.

The elements of an effective cyber security framework include:

- outlining the high-level approaches for managing information security. This needs to be supported by control standards to mitigate cyber security risks (for example, standards for implementing appropriate controls for network and application security, for managing software vulnerabilities, and for maintaining appropriate controls for physical security)
- having an information security policy, which defines an entity's objectives for managing information security. These objectives should align with the Queensland Government's *Information security policy (IS18:2018)*. An entity's policy should define the standards and frameworks that need to be implemented to achieve the objectives, such as the *Queensland Government Information Security Classification Framework*
- incorporating information security in an entity's enterprise governance. This could include having a security sub-committee of an entity's risk committee
- implementing mandatory information security training for all staff. This should cover:
  - an introduction to information security
  - maintaining and protecting passwords
  - device and equipment security
  - maintaining the security of confidential, personal, and sensitive information; email; and online security
  - information security incidents.

  This training can be delivered online to all users when they start work in the entity and should be mandatory for all users periodically throughout their employment/contract. Entities should also provide more targeted training for users who have access to sensitive information
- requiring periodic reporting (for example, monthly) from the entity's chief information officer, which should include the current security threat level, any identified vulnerabilities, any recent security incidents, email statistics (including the number of emails blocked for security reasons), any policy updates, and progress on cyber security risk awareness activities
- implementing a disciplinary process for users who commit information security breaches
- defining their information security risk appetite (the risk they're willing to take in order to meet their objectives) at the corporate and operational level
- allocating specific responsibility for coordinating a cyber risk assessment process for an entity's organisational units.

Managing key information assets

Information asset classification

DEFINITION

Information asset classification. According to the *Queensland Government Information Security Classification Framework*, entities should classify information assets according to business impact and implement appropriate controls according to the classification. The classifications are official, sensitive, and protected. (For more information, see Chapter 1.)

All three entities prioritise their information assets based on business criticality and value, but none of them had effectively implemented a process for security classification. Therefore, we could not be assured that they had effectively identified all the information assets that were at risk, or considered what controls were necessary to protect those assets.

We did note that all three entities recognised the importance of a security classification process and had plans to implement one.

At one of the audited entities, some business units advised us that the entity's central information technology services area provided insufficient guidance on how to implement information security classification on their sensitive information. The entity recently drafted an information security classification procedure and was preparing to implement this by:

- reviewing its information assets against confidentiality, integrity, and availability requirements, and in terms of whether they contain information that could be used to identify a person
- identifying information stewards (those who are accountable for information in their area of responsibility) and custodians (those who are responsible for ensuring the controls for managing information are applied on behalf of the steward).

For the other two entities:

- One had started to classify its information assets but had not yet classified 128 of its 269 applications. The entity has not yet used the 2018 *Queensland Government Security Information Classification Framework* definitions for security classifications. (It used the 2013 definitions.) It does, however, have plans to review and update its security classifications for its business applications as part of its implementation of a new information security management system.
- One did not have a policy and procedure for classifying its information assets, so it hadn't classified any of its information assets. It has plans to develop these now that it has employed a dedicated resource to manage cyber security.

Information and communication technology asset management

None of the three entities we audited was effectively managing its information and communication technology (ICT) assets to ensure it had a full record of the ICT assets it allocates to employees. In particular, the entities' processes for managing employee separations (when employees leave the entity) were not robust enough to ensure all ICT assets are returned. If assets are not returned, there is a risk these assets could still be used to access an entities' sensitive information.

One of the entities recently drafted an information management framework (policy and procedures, including security classification procedures). We tested this entity's asset register and found:

- there were 3,440 ICT assets (mostly laptops) without a custodian (out of 13,858 registered ICT assets—not all its ICT assets are registered, because the entity gives its organisational units discretion to register ICT assets)
- there were at least 581 instances where the custodians registered in the entity's asset register were no longer working for the entity.

For the other two entities:

- One records its ICT assets in a dedicated register but does not update the asset owners in a timely manner. At least 168 ICT assets (including portable and attractive items like desktops, laptops, tablets, and software) were recorded as being assigned to 118 employees who no longer work for the entity.
- One does not have adequate processes in place to identify and maintain an inventory of information assets.

QAO insight statement 2 shows the main elements for having an effective framework for managing information technology assets entrusted to employees.

QAO insight statement 2

Managing ICT assets

Managing ICT assets effectively is important in ensuring that unauthorised users do not have access to the entity's ICT devices to connect to the entity's network. Employees (and contractors) who do not return ICT assets when their employment with the entity ends could use their device for unauthorised purposes.

We observed the following practices are essential for managing ICT assets:

- maintaining a record of all ICT assets. If an entity grants discretion to its organisational units to register ICT assets, the entity does not have full awareness of all assets connecting to its network or of who the custodians of those assets are
- having a monitoring process to ensure information asset registers have been reviewed and kept up to date
- having a process to ensure employees return any ICT assets they have been accountable for when they cease their employment with the entity.

Entities can also consider using software solutions to manage and track their ICT assets, making it possible for them to disable access to a mobile device that a former employee still has access to.

Identifying and assessing cyber security risks

Methodology for identifying cyber security risks

Only one of the three entities could demonstrate that it periodically assesses its exposure to cyber security risks at an enterprise level and for business-critical applications. It has adequate risk assessment processes in place to identify information security risks posed to business-critical information.

This entity had established and documented an enterprise risk assessment process that covers all types of risks (including cyber security risks).

For the other two entities:

- One has cyber risk assessments performed in silos—or not at all in some areas—but to address this it has standardised its cyber security risk management processes across its many organisational units.
- One does not have a risk assessment process to identify and assess cyber security risks. While it has an enterprise risk management framework, the unit responsible for managing information security was not using the framework to develop a process for identifying and assessing cyber security risks.

While the entity without an assessment process registered some general cyber security risks in its risk registers, they were described at a corporate level, and not for specific information assets. This means the entity's risk treatments may not be enough to address risks to specific information assets.

Risk appetite

Risk appetite is the amount and type of risk that an organisation is willing to take in order to meet its strategic objectives. The absence of a risk appetite statement could lead to inconsistent implementation of risk mitigation processes.

Two of the three entities had risk appetite statements for cyber security. However, one of these entities had two different risk appetite statements for different organisational units. Because there is no distinction in the ICT services provided to these units, this could result in inconsistent direction on cyber security risk management.

Threat intelligence

Threat intelligence services can provide entities with valuable information on potential vulnerabilities that could affect their security posture (their ability to defend from and react to cyber attacks.) Entities can use this to identify cyber security risks.

One of the entities has implemented a formal process to actively receive, collate, and assess intelligence on threats and vulnerabilities in order to respond to emerging cyber security risks. It uses a variety of sources to gain threat intelligence data, including CITEC's threat intelligence and monitoring advisory service, AusCERT, and the Queensland Government Chief Information Office's vulnerability scanning service.

One entity conducts analytics of the deep and dark webs (where there is content not available on usual internet sites) to identify breached and exposed user credentials and passwords to help in identifying cyber risks. The other entity is subscribed to various services that alert it to current threats and vulnerabilities, but it has no formal process for analysing these alerts against its own risk criteria or taking appropriate action.

Security testing

Entities often conduct security testing—for example, information systems audits and penetration tests (simulated cyberattacks to evaluate the security of a system)—which is an effective means for helping to identify risks it may not have found through a desktop analysis. But when this happens, entities need to incorporate these results into their risk assessment processes so the risks can be monitored and tracked.

One of the entities recently engaged external experts to perform various cyber security reviews. They identified and assessed information security risks based on threats, vulnerabilities, and likelihood and impact of cyber security incidents.

However, we found the entity did not incorporate the risks, implications, and expected controls identified through these assessments into its cyber security risk register. Consequently, it did not implement adequate processes to track and manage the identified threats and risks.

We discuss the types of risks that entities can identify through threat intelligence and security testing in Chapter 4.

QAO insight statement 3 shows some of the important practices we observed for identifying cyber security risks.

QAO insight statement 3

Identifying cyber security risks

Identifying cyber security risks is important in ensuring an entity is aware of its risk exposure and assessing whether it has the right controls in place to mitigate those risks.

But entities must not only examine cyber security at an enterprise level (for example, 'we could be compromised'). They must also seek to understand how their information assets could be exposed to cyber risks (for example, application X is vulnerable because of Y and this could result in loss of sensitive data and reputational damage).

Entities should consider doing the following as a part of a process for identifying cyber security risks:

- identifying and classifying their information assets (the most critical information assets are also known as an entity's 'crown jewels'). Without this, they may not focus their cyber security risk mitigation activities on protecting their most important information assets
- defining their risk appetite for cyber security risks
- developing risk assessment processes for cyber security risks that integrate with their enterprise risk management framework. This should also include identifying any security risks associated with suppliers who provide ICT services (We discuss this further in Chapter 3.)
- identifying and assessing the exposure of their specific information assets to cyber security risks
- using threat intelligence services and security testing to inform their identification and assessment of cyber security risks. The Queensland Government Chief Information Office has developed a *Vulnerability management guideline* to assist entities in obtaining information on, evaluating, and acting on technical vulnerabilities within their environments
- testing their physical security, and the ability of attackers to access information from social networks, as part of security testing.

3. Mitigating cyber security risks

This chapter is about how effectively the three entities we audited implement cyber security risk mitigation strategies.

Introduction

To assess whether the three in-scope entities are effectively mitigating their cyber security risks, we:

- tested whether they have implemented the ‘Top 4’ of the 'Essential Eight' mitigation strategies published by the Australian Cyber Security Centre (ACSC) to help organisations protect their systems against cyber threats. The Top 4 strategies include application whitelisting, patching applications and operating systems, and restricting administrative privileges
- assessed whether they maintained formal risk management practices in their supply chains (for example, when they engage third parties to deliver information communication and technology (ICT) services). ICT managed service providers are a known target for attackers, because they require remote access to their customers' systems to deliver their services
- assessed whether they have implemented cyber security awareness programs to make general staff—and staff with privileged access (users who have administrative access to systems)—aware of their responsibilities for managing cyber security risks. The BDO Australia 2018–19 *Cyber Security Survey* states:

> Our trend data from survey results since 2016 outlines a consistent rise in phishing incidents through to 2018. In fact, it remains the most common incident experienced. Adversaries continue to target the human psyche, our inquisitiveness and general position of trust. Humans are continuing to prove to be a weak link in the layers of defence.

Application whitelisting

DEFINITION

Application whitelisting is a security approach for preventing the execution and spread of malicious code, and the installation or use of unauthorised applications. It ensures that only authorised applications can be run and installed.

Figure 3A shows our assessment of the three in-scope entities' implementation of whitelisting strategies.

Figure 3A
Application whitelisting

Processes and controls	Entity #1	Entity #2	Entity #3
Application whitelisting strategy and controls	🔴	🟡	🔴
Exception logs for failed execution of authorised codes, files, and programs	🟡	🟡	🟡
Restriction of dynamic link libraries (files that contain instructions that other programs can use), scripts, and installers	🟡	🟡	🟡
Application whitelisting using methods approved by the Australian Signals Directorate (cryptographic hash rules, publisher certificate rules, and path rules)	🔴	🟡	🔴

Legend: 🟢 Process/control implemented and operating effectively 🟡 Control partly implemented or evidence of some compensating controls 🔴 Control not implemented and compensating controls ineffective or lacking.

Source: Queensland Audit Office.

None of the three entities had a strategy for implementing application whitelisting, but one of the entities has started to implement application whitelisting in its server and desktop environment. All three entities had some compensating controls in place, but we found these were limited in their effectiveness in fully mitigating the risk that a whitelisting strategy is meant to address.

We observed two challenges the audited entities face in implementing application whitelisting:

- Entities that are tasked with collaborating and experimenting in their ICT environment find it challenging to implement application whitelisting. This is because, while application whitelisting makes an entity's systems more secure, it also limits what users can run on the network.

  One of the three entities had not implemented application whitelisting, because it was concerned about making sure it provided its users with a collaborative working environment, while at the same time protecting sensitive information. Entities in this position should still document a whitelisting strategy to show how they plan to address the risk of malicious code being executed in their environment. This could include implementing application whitelisting on devices used by users who have access to sensitive information.

- Implementing application whitelisting is complex, because entities need to account for all the different layers of technology in their networks (such as desktops and servers, and in some cases, different versions of desktops and servers). Otherwise, they can be left exposed on parts of their network that are not protected. One of the entities implemented an application whitelisting solution within its Windows 10 desktops and newer servers, but we found:
  - 93 per cent of its servers hosting applications did not have a whitelisting solution because most of the entity's application hosting servers use old versions of servers or servers reaching end of life. We recognise the entity does not want to implement an application whitelisting solution on servers it plans to replace soon, but in the interim, there is a risk that unauthorised/malicious executables could be run and installed on these servers.
  - 33 per cent of the desktop/laptop computers did not have a whitelisting solution, because they used an older version of the Windows operating system. The entity has a plan to upgrade (by October 2019) all those devices to the most current Windows operating system version (Windows 10), which is the operating system for which the entity has implemented application whitelisting.
  - The entity allowed Windows 10 desktops to run scripts based on their user access privileges, which means privileged users can run scripts from untrusted sources. The entity relies on the user signing an undertaking to use their privileges for the approved purposes. However, this is not an effective compensating control, because a potential attacker could exploit this weakness if they successfully compromised a privileged user account.

QAO insight statement 4 describes how entities can implement a whitelisting strategy, based on guidance material from the Australian Signals Directorate and our observations in this audit.

QAO insight statement 4

Application whitelisting

Application whitelisting is a control to prevent unauthorised and malicious software being installed on an entity's network, allowing an attacker to compromise the network and gain access to data.

Entities should consider:

- developing a whitelisting strategy and implementing a solution on their desktops and servers to prevent the unauthorised installation of codes, files, software, and programs

  Having a documented whitelisting strategy demonstrates than an entity has considered the appropriate level of controls to implement. For example: overly restricting whitelisting processes creates unnecessary costs, whereas less restricting whitelisting creates a security risk. A whitelisting strategy will help entities ensure they have the necessary processes in place

- establishing a process for deciding and monitoring which applications should be whitelisted, and a process that enables users to formally request approval for applications to be whitelisted
- capturing user exception logs to identify failed attempts to execute unauthorised codes, files, software and programs
- using methods recommended by the Australian Signals Directorate—cryptographic hash rules, publisher certificate rules (combining both publisher names and product names), and path rules (ensuring file system permissions are configured to prevent unauthorised modification of folder and file permissions, folder contents and individual files)
- If entities use older versions of operating systems and decide not to implement application whitelisting on those desktops and servers until they upgrade them, their whitelisting strategy should identify how they will address this risk in the interim.

If entities have concerns about implementing a whitelisting solution in their entire networks, they should consider the following advice from the Australian Signals Directorate guide on application whitelisting:

> Implementing application whitelisting across an entire organisation can be a daunting undertaking; however, implementation on at least workstations of high-risk users such as senior managers and their staff; system administrators; and staff members from human resources, sales, marketing, finance and legal areas can be a valuable first step.

Patching operating systems and applications

DEFINITION

Patches are released by software and hardware vendors to fix known vulnerabilities that attackers could exploit (as well as to address a software flaw or to improve the stability of an application).

Figure 3B shows our assessment of the three in-scope entities' implementation of patch management strategies.

Figure 3B
Patching operating systems and applications

Processes and controls	Entity #1	Entity #2	Entity #3
Patch management strategy	🟡	🟢	🔴
Patching approach and processes	🟡	🟢	🟡
Patching and mitigating extreme risk security vulnerabilities	🟡	🟢	🔴
Patching and mitigating below extreme risk security vulnerabilities	🟡	🟢	🔴
Replacing/updating legacy (outdated) systems to vendor-supported versions	🟡	🟢	🔴
Mitigating vulnerability risks when patches are not available	🟢	🟢	🔴

Legend: 🟢 Process/control implemented and operating effectively 🟡 Control partly implemented or evidence of some compensating controls 🔴 Control not implemented and compensating controls ineffective or lacking.

Source: Queensland Audit Office.

One of the audited entities demonstrated effective controls for patch management. We observed that this entity:

- has a strategy for patch management. It provides its security governance group with periodic reports on its patching activities
- has mature processes for implementing patches. It uses trusted sources to ensure the integrity and availability of patches required for its environment
- effectively implemented a patch management process to ensure it gave priority to implementing patches assessed as extreme risk and had adequate processes for implementing patches assessed as below extreme risk. Its policy meets the Australian Signals Directorate's recommendation for patching extreme risk vulnerabilities within 48 hours (it aims for 24 hours) and for patching vulnerabilities assessed as below risk as soon as possible
- updates or replaces any operating systems, applications, and hardware devices that are no longer supported by their vendors to vendor-supported versions or alternative vendor-supported versions. It has a procedure for maintaining vendor support for all ICT assets
- implements other mitigating approaches when patches are not available to address security vulnerabilities (for example, when an extreme risk event like the 'wannacry' virus occurs, and patches are not readily available or are difficult to deploy with a high degree of success). The entity escalates these issues to its security governance committee, where such non-compliances are risk assessed, addressed, and tracked.

Case study 1 shows how this entity assesses its risk exposure in relation to security vulnerabilities that can be addressed by patches released by software vendors.

Case study 1

Assessing the risk ratings of vulnerabilities that can be addressed by vendor-released patches

This entity has processes in place for identifying, managing, and assessing known vulnerabilities in security that are relevant to its operating environment. It does this in terms of how critical and how severe the vulnerabilities could be. To assess these risks, the entity:

- conducts its own monthly vulnerability scanning
- assesses results from the Queensland Government Chief Information Office's (QGCIO) monthly managed vulnerability scanning service
- assesses results from AusCERT's security bulletin to identify vulnerabilities relevant to its standard operating environment
- assesses QGCIO's *ICT Alert* emails for vulnerabilities currently being actively exploited by threat actors (that is, people or entities acting maliciously) and applies appropriate risk treatments (including patch management) to address identified cyber risks.

The entity reports these activities to its information security governance committee on a monthly basis.

The entity also actively participates in the QGCIO cyber security Community of Practice meeting, where they discuss and share ideas with QGCIO, CITEC and other departmental vulnerability practitioners and security experts.

These activities demonstrate that the entity has a risk-managed approach for identifying, managing, and addressing cyber security risks in relation to vendor-released patches.

Source: Queensland Audit Office.

The other two entities did not have mature processes for patch management. We observed:

- One entity has an established process for patch management, applies these processes to workstations, and sources the patches appropriately. But it applies patches on servers on an ad hoc basis. It has appropriate policy settings in place for mitigating extreme risk security vulnerabilities, but it needs a more robust risk assessment process to enable this. It implements other mitigating approaches when patches are not available to address known security vulnerabilities.
- One entity does not have a strategy or operating procedure for patch management. Its approach for patching systems is ad hoc and mostly left up to the user. It does not have a specific timeframe for implementing patches, regardless of the risk severity the vendor patches are designed to address.

  It still uses many unsupported applications and operating systems and it does not have a process to manage those legacy systems that it still needs. When patches are not available for known security vulnerabilities, it relies on staff being prudent and knowledgeable in identifying and reporting these scenarios to management to determine mitigation strategies. It does not have a procedure for undertaking risk assessments of unsupported applications and operating systems. As a result, it cannot identify the level of cyber security risk and develop appropriate resolution, prevention, containment, and detection strategies.

QAO insight statement 5 describes how to effectively implement a patching strategy, based on guidance material from the Australian Signals Directorate and our observations in this audit.

QAO insight statement 5

Patching operating systems and applications

Because patches indicate that there are vulnerabilities in software that attackers could exploit, it is important that entities implement them in a timely manner, using a risk assessment process. Otherwise, they could leave themselves unnecessarily exposed to potential attacks.

Effective controls over patch management include:

- patching vulnerabilities in operating systems, applications, drivers, and hardware devices within the following timeframes recommended by the Australian Signals Directorate:
  - 48 hours for extreme risk vulnerabilities
  - as soon as possible (for example, within the next patch cycle) for other vulnerabilities
- having a risk assessment process for vendor-released patches to ensure the entity applies patches within the timeframes specified in their patch management strategy or operational procedures. To make this process more efficient and effective, entities should use vulnerability scanning services, such as those provided by the QGCIO. These services identify when vulnerabilities have already been exploited. If a vulnerability has already been exploited, it increases the risk that an entity could be attacked
- having an effective governance process to provide senior management with information they need on what vulnerabilities the entity has identified and what progress it has made in patching them
- updating or replacing any operating systems, applications, and hardware devices that are no longer supported by their vendors, to a vendor-supported version or alternative vendor-supported version
- implementing mitigating approaches for security vulnerabilities when patches are not available
- implementing regular patching for all systems that are externally accessible.

Restricting administrative privileges

The Australian Signals Directorate guidance on restricting administrative privileges states:

> Users with administrative privileges for operating systems and applications are able to make significant changes to their configuration and operation, bypass critical security settings and access sensitive information. Domain administrators have similar abilities for an entire network domain, which usually includes all of the workstations and servers on the network.
>
> Adversaries often use malicious code (also known as malware) to exploit security vulnerabilities in workstations and servers. Restricting administrative privileges makes it more difficult for an adversary's malicious code to elevate its privileges, spread to other hosts, hide its existence, persist after reboot, obtain sensitive information or resist removal efforts.
>
> An environment where administrative privileges are restricted is more stable, predictable, and easier to administer and support, as fewer users can make significant changes to their operating environment, either intentionally or unintentionally.

Figure 3C shows our assessment of the three in-scope entities' implementation of strategies to restrict administrative privileges.

Figure 3C
Restricting administrative privileges

Processes and controls	Entity #1	Entity #2	Entity #3
Managing privileged accounts	🟡	🟢	🟡
Restricting database administrator access	🟢	🟢	🟢
Having secure communication for remote system administration tasks	🟡	🔴	🔴
Restricting internal and email access on privileged accounts	🟡	🟢	🟡
Logging and monitoring privileged operations	🟡	🟢	🔴

Legend: 🟢 Process/control implemented and operating effectively 🟡 Control partly implemented or evidence of some compensating controls 🔴 Control not implemented and compensating controls ineffective or lacking.

Source: Queensland Audit Office.

Managing privileged accounts and restricting database administration access

One of the three entities we audited had implemented effective controls for managing privileged user access. It had processes for assigning and reviewing privileged user access, and it automatically revokes access after 12 months, after which the user must reapply. It only provides privileged user access to authorised business system administrators or central ICT support staff, under an agreement with the business unit.

For the other two entities:

- One has designed and implemented a documented process for administering and managing privileged user access, However, some business units that operate their ICT environments independently do not apply a robust process for managing administrator privileges. We observed some of the business units provided local administrator access (which is privileged access) to all users.
- The other has documented responsibilities for managing user access, but it has not documented the process for administering, managing, monitoring, and reviewing privileged user access. While the entity restricts and manages privileged user access through role-based access and through separate accounts for standard and privileged access, these restrictions are not fully effective because
  - there is evidence of privilege-creep. One user still had domain administrator access that we confirmed was no longer required for their role
  - the entity does not periodically monitor or review privileged user access. It does not have a process that describes how often administrative privileges should be reviewed.

We found all three entities implemented effective controls to minimise and restrict database administration privileges.

Having secure communication for remote system administration tasks

DEFINITION

A **jump box** (or server) provides a control to ensure access to a secure server cannot be obtained from a less secure zone on the corporate network. For example, a jump box ensures that if a user needs to access an administrative account from outside the secure zone, it can only do so through the jump box and not directly from a normal user zone.

None of the three entities we audited had fully implemented controls to ensure administration tasks can only be performed through a secure connection (like a jump box or a virtual private network (an encrypted connection from a device to the network)). This creates a risk that if an attacker gains access to their networks and compromises the password for an administrator account, they could access the account without requiring any further authentication beyond the initial compromised password.

Restricting internet and email access on privileged accounts

Privileged user accounts should be restricted from business-as-usual activities such as reading email and web browsing. This is to reduce the risk of a privileged account being used to download malicious software or being subject to a phishing attack.

Only one of the three entities had implemented controls to prevent privileged user accounts from being used for reading emails, opening attachments, and browsing the web or obtaining files via internet services such as instant messaging or social media. Privileged accounts have no email or internet access.

For the other two entities:

- One does not prevent privileged user accounts from reading emails, opening attachments, browsing the web, or obtaining files via internet services such as instant messaging or social media. It has, however, implemented some compensating controls.
- The other does not provide email for privileged user accounts, but administrators could set up their own email accounts. In addition, privileged accounts have unrestricted access to the internet (which includes webmail).

Logging and monitoring of privileged operations

We observed varied results across the three entities with regards to logging and monitoring of privileged accounts.

One entity had implemented effective controls to log and monitor privileged user activity. It has designed and implemented processes and controls to securely maintain, monitor, and review audit logs for privileged user account activities for all critical ICT systems and applications. It logs all administrator activities and performs various monitoring and alerting activities.

It uses a security event management system (which provides real-time analysis of security alerts generated by network devices and software) to collect and analyse a minimal amount of audit logs, but it has plans to increase the collection of logs from different sources to better detect anomalous events.

For the other two entities:

- One has implemented logging of privileged user account activities on the operating systems we tested, and it keeps these logs for a moderate period. However, it does not monitor or receive alerts of privileged user account activities.
- The other has no strategy or policy for securely maintaining, monitoring, or reviewing audit logs for privileged user account activities for its critical ICT systems and applications. It has implemented limited logs of activities in general, but they do not include monitoring or alerting and do not target privileged activities. In addition, even these limited logs are overwritten, on average, within 24 hours.

QAO insight statement 6 describes how entities should implement a process to restrict administrative privileges, based on guidance material from the Australian Signals Directorate and our observations in this audit.

QAO insight statement 6

Restricting administrative privileges

If an attacker successfully compromises privileged user accounts, they can use these accounts to gain full access to an entity's systems and information.

Effective controls for restricting administrative privileges include:

- reviewing and monitoring privileged user access on a periodic basis
- logging and monitoring privileged user account activities for all critical ICT systems and applications
- requiring remote system administration to be only performed through a secure connection
- restricting internet and email access on privileged accounts.

Managing security risks in the supply chain

All three entities we audited need to improve their practices for managing security risks in their supply chain. Two of the three entities had not yet established formal processes for managing security risks associated with suppliers who provide ICT services. The other had defined the ICT procurement process and the roles and responsibilities for the general procurement process, but it had not defined the information security-related roles and responsibilities for managing information security risks in its supply chain.

We found an example of good practice at one of the entities in relation to its formal risk management practices for its major enterprise resource planning software solution. It routinely assesses if the supplier is meeting its contractual obligations. This includes conducting regular audits and other forms of checks and balances such as daily monitoring activities, and obtaining auditor reports of the supplier from a third party to confirm they are meeting contractual obligations.

One of the entities was still developing its process for managing third-party risks in the supply chain. Because these administrative process documents were not finalised and operational, we could not be assured that the entity:

- consistently manages and monitors cyber security risks in its supply chain
- manages and monitors security-related service-level key performance indicators against contractual terms and conditions.

We also noted:

- it did not conduct regular review and monitoring of third-party ICT supplier services to ensure they maintain security levels in alignment with its own relevant information/cyber security policies and procedures (or international security best practices)
- it did not have a risk assessment process for its third-party suppliers to determine the suitability of potential external service providers from an information security perspective.

One of the other entities did not include specific information security provisions in its contracts with third-party providers. It relied on standard contract clauses. We reviewed its hosted services agreement and found there were no information security provisions. While there were references to operational practices such as patching, there were no provisions for information security requirements with which the vendor needed to comply. This means there was no process for monitoring, reviewing, or auditing the vendor's compliance with information security requirements.

We found one of the entities had adequate processes for ICT contracts that are valued at more than $10,000. The processes aim to provide clarity and accountability of security responsibilities throughout the life of the relationship. But the entity did not have this in place for contracts under $10,000.

We found an instance where a business unit procured a small ICT solution without obtaining endorsement from the central ICT area. The system was later found to not meet the entity's security criteria and was incompatible with its infrastructure. An ICT solution, irrespective of dollar value, can introduce risks that need to be identified and managed.

QAO insight statement 7 shows some of the important practices for a third-party risk management process.

QAO insight statement 7

Managing security risks in the supply chain

Entities need to understand the risks of engaging third parties to deliver ICT services. If third parties do not understand and comply with the security requirements of the entity, the entity can be exposed to cyber risks it cannot effectively control.

Entities need to establish effective risk management practices for third parties that cover the initial engagement of the supplier and the ongoing relationship with the supplier. As entities increase their use of cloud services to deliver software solutions, this becomes even more important.

An effective process for managing third-party risks in the supply chain includes:

- having a risk assessment process to determine the suitability of potential external suppliers
- defining information security responsibilities with which suppliers must comply. Entities should not rely on standard contract clauses. Instead, they should be specific about what security expectations they have of the supplier
- having processes for starting and finishing engagements with external suppliers
- regularly monitoring, reviewing, auditing, or evaluating service delivery to ensure suppliers are meeting their security obligations.

Cyber security awareness training

Two of the three in-scope entities operated cyber security awareness programs to enhance staff awareness of their roles and responsibilities for protecting against cyber security risks. Both entities provided cyber security awareness training to all staff, and one of these entities has provided more targeted training to users who have access to sensitive data.

We found the third entity did not provide any cyber security awareness program. This entity is therefore more susceptible to the type of attacks that take advantage of weaknesses in human controls, like phishing attacks.

DEFINITION

Phishing is a method attackers use to gather personal information from users, using deceptive emails and websites. The attackers masquerade as a trusted person or business the person is used to doing business with, in order to get the user to release sensitive information or to download malicious software onto their computer (malware) that allows the attacker to take control of the computer and obtain data.

One of the entities conducts simulated phishing campaigns against randomly selected employees to raise staff awareness and vigilance. It monitors the percentage of staff who fall for the phishing campaigns to determine the effectiveness of its cyber security awareness programs.

QAO insight statement 8 shows key elements in an effective cyber security awareness training program.

QAO insight statement 8

Cyber security awareness training

While entities can have strong governance processes and technical controls to mitigate cyber security risks, people can be the weak link in the layers of defence.

Providing an effective cyber security awareness program includes:

- developing and implementing mandatory cyber security awareness training for all staff, to be completed during induction and at regular periods during employment
- delivering targeted training to higher-risk user groups, such as senior management, staff who have access to sensitive data, software developers, system administrators, and third-party providers
- recording and monitoring whether all staff have completed their required cyber security awareness training
- conducting campaigns to test the adequacy of staff vigilance to risks such as phishing and tailgating (following a person into an office), so entities can assess and improve their awareness programs.

4. Testing cyber security controls

This chapter is about how effectively the three entities implement cyber security controls to prevent access to their sensitive data.

Introduction

We commissioned security consultants to perform an open source threat intelligence assessment and a red team security assessment for the three entities within the scope of this audit. We used these assessments to identify vulnerabilities that entities need to address to protect their sensitive data from attacks.

DEFINITION

Open source threat intelligence assessment

This involved determining whether any sensitive information about the three entities could be obtained from public sources.

Information about entities that can be sourced publicly can be used by a potential attacker to determine avenues for gaining access to entities' systems.

Our security consultants also investigated publicly available information from the internet and the hidden web (also known as the 'deep web' and 'dark web', this has content not available on usual internet sites). They assessed numerous leaked email address lists and sources, as well as online applications used to generate targeted phishing email address lists. They also scrutinised various social networking sites and organisational websites.

Red team security assessment

A red team engagement tries to find the quickest method to access an entity's security mechanisms and compromise its sensitive applications and data. In doing so, it considers the target and resources available, and may attempt social engineering, physical entry, and data exploitation. The objective of this assessment was to simulate a targeted attack against the three entities to gain unauthorised access to their sensitive applications and data.

We note that a red team assessment does not include a thorough assessment of the network and physical controls, but only of what is required to access the specified target. This means, for example, that if the team successfully accessed the target by penetrating the external network, it should not be assumed this means the physical security controls are adequate. It simply means that it was easier to gain access through external network penetration.

Overview of results from security testing

Red team security assessment

For each of the three entities, we nominated a narrow set of targets for our security consultants to test, based on our understanding of the entities' key information assets.

In all three instances, our security consultants were successful in compromising at least one target we set for each entity. While the attack path our security consultants used differed across the three entities, based on what they identified as the easiest path to attack, there were some common issues that made their task easier.

Open source threat intelligence assessment

For two of the three entities we audited, our security consultants identified multiple services that exposed information an attacker with no inside knowledge of these two entities could use to build up a catalogue of information. We found that overall availability of sensitive information relating to the other entity was minimal.

In the following sections, we outline the main learnings for all entities from the results of our security testing.

Physical security controls

At one of the three entities we audited, our security consultants gained initial access to the network through poor physical security controls.

Our security consultants were not prompted for identification at any point when accessing facilities. It was possible to walk from the lifts, past the reception desk, and tailgate employees into the entity's offices. Upon accessing the office, our consultants were able to sit down at employee desks and connect a malicious device to the network.

This facilitated direct access to the entity's internal assets and increased the available ways to target the entity.

User account management

Password practices

At all three entities, we found easily guessable passwords made it easier for our consultants to compromise user accounts and use them to gain control of the entities' networks. At one entity, our consultants were able to crack and recover clear text passwords for over 6,000 user accounts. They cracked the majority of these in less than three minutes.

They were also able to extract passwords from configuration files (files that are used to structure the parameters and settings for some computer programs) and system memory. At another entity, our consultants were able to crack and recover clear text passwords for over 800 user accounts. Again, they were able to crack most of these accounts in less than three minutes.

Our consultants found users could set common account passwords typically found in password breaches (for example, Password1! and welcome1!). They also found instances where passwords contained:

- predictable permutations, given the frequency of change (like changing numbers on the end of a password)
- an application name
- dates
- the entity name or service.

Figure 4A shows the top nine common base passwords used in the accounts cracked in this assessment at two of the entities. The actual password includes the base word and may have additional characters, for example, Welcome01.

Figure 4A
Common base passwords

Entity X		Entity Y	
Base word	**% of cracked passwords**	**Base word**	**% of cracked passwords**
welcome	16.2	newuser	8.7
password	3.97	password	3.5
monday	1.58	pa55word	3.26
summer	0.86	Entity service	0.97
march	0.83	Entity name (1)	0.97
passw0rd	0.80	Entity name (2)	0.72
april	0.80	monday	0.72
p@assword	0.57	thursday	0.72
february	0.54	welcome	0.60

Source: Queensland Audit Office.

Known password breaches

Our consultants found over 500 user accounts, associated with the three entities' email addresses, to have passwords that have been compromised and disclosed in multiple data breaches that are publicly available. These passwords were associated with services such as Adobe, Dropbox, LinkedIn, and MySpace (and other unattributed breaches).

These user account and password leaks do not indicate the entities' accounts were, or could be, breached. However, a persistent attacker could find valid passwords that the entities' users reuse across multiple accounts.

Entities should make their staff aware of the risk they create for their entities when they use the same user account and passwords on multiple online services.

Accounts with administrative privileges

At two out of the three entities, our consultants found that users who had administrative privileges to computers in the entities' environments actively used the accounts for business-as-usual purposes. This made it easier for our consultants to compromise these accounts to gain control of the entities' systems.

Multi-factor authentication

For one of the entities, our consultants found users could remotely gain access to the entity's internal network without two-factor authentication. For two of the three entities, our consultants found staff and administrators were allowed access to sensitive internal servers without having to supply multi-factor authentication.

The combination of easily guessable passwords and the lack of two-factor authentication for:

- external-facing services (such as a website that enables a user to log in to an entity's service) could enable an attacker to gain access to the entity's network through password guessing
- internal services could enable an attacker who can gain access to a valid highly privileged username and password to use those login credentials to gain access to sensitive internal network servers.

Network segmentation

At all three entities, once our consultants could access the internal networks, they could move laterally within their networks to compromise additional systems. There were minimal security controls restricting access between key systems.

During the engagement, our consultants used several user credentials to move throughout the entities' networks to target servers on adjacent network segments and within the same network segments.

If a rogue employee or attacker compromised a single server within these networks, they would be able to use the access or credentials to target other servers.

Outdated systems

At two of the entities, our consultants identified numerous systems were running outdated applications and operating systems (such as Windows XP and Server 2003) that had not been supported by the vendor for several years. This infrastructure has several known critical vulnerabilities that allows unauthenticated users the ability to execute arbitrary code remotely, granting full system access, and the ability to further attack internal services from trusted locations. This malicious code is publicly available and distributed as default packages within commonly used hacking tools.

Microsoft ended support for Windows 2003 on 14 July 2015 and Windows XP on 8 April 2014. Software and operating systems that are out of support will no longer routinely receive new security patches. As a result, security vulnerabilities are likely to be present and unlikely to be fixed by vendors.

Descriptive subdomains

DEFINITION

A **subdomain** is an internet domain which is part of a primary domain. For example, a primary domain may be xxx.qld.gov.au and a subdomain of this could be yyy.xxx.qld.gov.au.

For all three entities, our consultants identified some subdomains that may give attackers insight into their ICT environment. We observed that:

- one entity's naming of its subdomains could provide attackers with insight into the online services it uses. As an example of how this information could be useful to an attacker, when a new vulnerability about a service is published, the attacker will scan the internet for subdomains that use that service
- two of the entities' naming of their subdomains indicated which subdomains are non-production environments (for example, development.entityname.qld.gov.au), which may not be as strongly secured as production environments, thus offering easier targets to attack.

Insecure encryption channels

All three entities used online application hosts (websites that allow users to use a software application over the web, for example, a mapping service) that do not use encryption. This weakens the online application’s integrity and could allow an attacker to manipulate communication between users and online services.

Security monitoring and response

DEFINITION

Endpoint protection is security software that protects end-user devices, such as desktop computers, laptops and mobile devices.

We found none of the three entities detected our security testing or prevented our consultants from accessing our set targets. In all three cases, we needed to advise the chief information officers that our security consultants had succeeded in accessing the targets.

This indicates that, should a malicious actor gain access to the internal networks of these entities, their monitoring systems may not be alerted. Once the servers were compromised, there were no instances during testing where any endpoint protection was identified to prevent further progression within the entities' networks.

Appendices

A. Full responses from entities

As mandated in Section 64 of the *Auditor-General Act 2009*, the Queensland Audit Office gave a copy of this report and an opportunity to comment to the relevant entities.

This appendix contains the formal response we received.

The head of this organisation is responsible for the accuracy, fairness, and balance of their comments.

Comments received from Queensland Government Chief Information Officer, Queensland Government Chief Information Office

OFFICIAL

Queensland Government

Department of
Housing and Public Works

For reply please quote: *QAO 9184P*

13 September 2019

Brendan Worrall
Queensland Auditor-General
qao@qao.qld.gov.au

Dear Mr Worrall

Performance audit on managing cyber security risks

Thank you for your letter of 9 August 2019 enclosing the draft *Performance audit on managing cyber security risks*. We welcome the opportunity to provide our views on the draft.

QGCIO have not identified significant gaps in the draft QAO findings. Although some further description of whether the entities audited were departments, authorities, boards or commissions would provide additional perspective and may avoid some misinterpretation and assumptions by the audience.

In QGCIO's experience, the entity size and whether they are departments have a correlation with increased cyber security maturity. Smaller government bodies, like authorities, boards and commissions are observed by QGCIO to generally have lower cyber security maturity.

To date there has been also an unclear authority for Queensland Government statutory bodies to meet the requirements of the QGEA. I note that, following 18 months of close collaboration between Queensland Treasury and my office, the scope and authority of the QGEA policy framework has been clarified somewhat in updated Financial Performance and Management Standard (FMPS 2019). We also note that, the Queensland Local Government Act 2009 does not bind councils to follow the QGEA or the FPMS.

Our view is that agencies are improving but show the effects of a historic lack of focus on cyber security. Cyber security processes have improved in government, but generally agencies would benefit from further enhancement of maturity in all areas of risk management, not least cyber security.

QGCIO's program of centre-led cyber security has seen significant improvements over the past 5 years. Whole of Government activities, such as inter departmental cyber exercises, that have been conducted as recently as March 2019, have had positive effects on agency cyber capability. Whole of Government services, such as the vulnerability scanning service

Queensland Government Chief Information Office
Level 24 111 George Street Brisbane QLD 4000
PO Box 15086 City East QLD 4002 Australia

Telephone +617 3215 3900
Website www.qgcio.qld.gov.au
ABN 41 841 375 926

OFFICIAL

and phishing simulation platform are also enjoying significant take-up. However, the centrally funded cybersecurity program is not currently focused to provide services beyond departments.

QGCIO will consider the audit findings and use it to inform its program and advice to agencies over the next three years.

Should you or your Office require any further information, please contact Mr Robert Champion, a/Queensland Government CISO

Yours sincerely

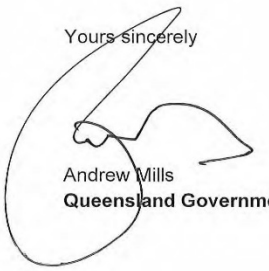

Andrew Mills
Queensland Government Chief Information Officer

B. Performance engagement

This audit has been performed in accordance with the *Standard on Assurance Engagements ASAE 3500 Performance Engagements*, issued by the Auditing and Assurance Standards Board. This standard establishes mandatory requirements and provides explanatory guidance for undertaking and reporting on performance engagements.

The conclusions in our report provide reasonable assurance that the objectives of our audit have been achieved. Our objectives and criteria are set out below.

Audit objective and scope

This audit examined whether entities effectively manage their cyber security risks.

It addressed this by assessing whether entities:

- understand and assess to what extent their information assets and organisational processes are exposed to cyber security risks
- design and implement effective information controls to mitigate identified cyber security risks.

Scope exclusions

We did not, as part of this audit, examine the effectiveness of activities conducted by the Queensland Government Chief Information Office. Only one of the three entities in this audit use the services of the Queensland Government Chief Information Office.

Entities subject to this audit

We selected three entities for this audit. We have not named the three entities in this report as we do not want to compromise their security by publicly identifying their security vulnerabilities.

We provided each of the in-scope entities with a detailed report on the risks we identified through our detailed testing, specific to their entity.

We acknowledge the three entities have different levels of resourcing and capability for managing cyber security risks.

Audit approach

We conducted the audit in accordance with the Auditor-General of Queensland Auditing Standards—September 2012, which incorporate the requirements of standards issued by the Australian Auditing and Assurance Standards Board.

The audit included:

- interviews with staff from the three in-scope entities
- review of documents and analysis of data
- a red team assessment for each of the three entities. A red team engagement tries to find the quickest method to access an entity's security mechanisms and compromise its sensitive applications and data. In doing so, it considers the target and resources available, and may attempt social engineering, physical entry, and data exploitation
- an open source threat intelligence assessment to determine whether any sensitive information about the three entities could be obtained from public sources
- testing whether the entities had implemented the 'Top 4' of the 'Essential Eight' mitigation strategies published by the Australian Cyber Security Centre to help organisations protect their systems against cyber threats
- interviews with staff from the Queensland Government Chief Information Office (as a stakeholder).

C. Information security policy (IS18:2018)

Figure C1 shows the policy requirements of *Information Security Policy (IS18:2018)*.

Figure C1
Information security policy (IS18:2018)

Policy requirement	Description
Policy requirement 1: Agencies must implement an ISMS based on ISO 27001.	Agencies must implement and operate an information security management system (ISMS) based on the current version of ISO 27001 Information technology—Security techniques—Information security management systems—Requirements. The scope of the ISMS will include the protection of all information, application and technology assets.
Policy requirement 2: Agencies must apply a systematic and repeatable approach to risk management.	Risk management is an integral part of operating an information security management system where risks must be considered at a business level. Agencies must adopt a risk management framework by integrating their ISMS into their corporate risk management processes.
Policy requirement 3: Agencies must meet minimum security requirements.	To ensure a consistent security posture and promote information sharing, Queensland Government agencies must comply with the: • *Queensland Government Information Security Classification Framework* (QGISCF) • *Data encryption standard* • *Queensland Government Authentication Framework* (QGAF) • Australian Signals Directorate's 'Essential Eight' strategies.
Policy requirement 4: Agency accountable officers must obtain security assurance for systems.	Every system is unique and security assurance should be applied sensibly and appropriately. Accountable officers must obtain security assurance to establish an understanding of information security protections and adherence to information security policy. The level of security assurance applied to systems must be based on the criticality/significance of the system, using the business impact levels determination methodology outlined in the QGISCF.

Source: Information security policy (IS18:2018) developed by the Queensland Government Chief Information Office Cyber Security Unit.

D. Glossary

Term	Definition
Application whitelisting	A security approach to prevent the running and installation of malicious code and unauthorised applications. It only allows programs that have been explicitly approved to be run and installed.
Cyber security	A process for protecting an entity's information by preventing, detecting and responding to cyber attacks. Such attacks could be through breaches of physical and network security, and through using information obtained through social networks.
Database administrative privileges	Administrative access to database systems that enables a user to create user accounts, set passwords, and manage and maintain databases that can contain sensitive data.
Encryption	A process for encoding a message or file so that it can only be read by authorised people. It makes sensitive data more secure and reduces the likelihood that an unauthorised person could intercept it to view it.
Endpoint security	Involves making sure there is security on the endpoints (potential entry points) of a network, like laptops and wireless and mobile devices.
Essential Eight	The Australian Cyber Security Centre (ACSC) has compiled a list of mitigating strategies entities can use to improve their ability to protect against cyber security risks. It has developed eight mitigation strategies that it says should be implemented as a baseline where practicable. They are known as the 'Essential Eight.'
Entities	We use the term 'entities' in this report to refer broadly to all Queensland public sector entities (including departments and statutory bodies) and local governments.
Execute	The process of running a computer software or command.
Information asset	A collection of data that is recognised as having business value and enables an entity to perform its business functions.
Jump server	A jump box (or jump server) ensures that access to a secure server cannot be obtained from a less secure zone on the corporate network.
Malicious or threat actor	An individual, group of individuals, or entity that attempts to conduct malicious activities against an entity by taking advantage of vulnerabilities to gain unauthorised access to systems and data.
Network segmentation	Involves segregating part of a computer network. This helps to reduce what is available to an attacker if they successfully compromise part of the network.
Open source threat intelligence assessment	An assessment that involves investigating publicly available information from the internet and the hidden web to determine whether any sensitive information about an entity can be obtained from public (or 'open') sources. The hidden web includes the 'deep web', which is the part of the world wide web where content is not discoverable using standard search engines, and the 'dark web', which is the part of the world wide web only accessible using special software.
Patches	Released by software and hardware vendors to mitigate known vulnerabilities that attackers could exploit (as well as to address a software flaw or to improve the stability of an application/program).

Term	Definition
Phishing	A fraudulent scamming attempt to obtain sensitive information from an end user (for example, username, passwords, and credit card information). For example, asking a user to click on a link that results in malicious software being installed.
Privileged user access	Administrative access to systems. For example, a user with privileged user access can create user accounts and set passwords, configure systems, have access to sensitive data, and execute other software and scripts.
Red team assessment	A red team engagement tries to find the quickest method to access an entity's security mechanisms and compromise its sensitive applications and data. In doing so, it considers the target and resources available, and may attempt social engineering, physical entry, and data exploitation.
Security posture	The security status of an entity's networks, information, and systems based on its resources (for example, people, processes, and technology) and ability to defend the entity from cyber attacks and to react as the situation changes.
Server	A computer program or a device that is dedicated to managing network resources to provide services to computer programs on end-user devices (for example, desktops, laptops, phones, and tablets).
Subdomains	A subdomain is an internet domain that is part of a primary domain. For example, a primary domain may be xxx.qld.gov.au and a subdomain of this could be yyy.xxx.qld.gov.au.
Two-factor authentication (or multi-factor authentication)	Requires more than one authentication method to gain access to a system, for example, a username and password, plus a code sent to a mobile phone.
Virtual private network (VPN)	Provides additional security to protect sensitive data on a corporate network. It provides an encrypted connection from a device to the network over the internet. It allows the user to work remotely and prevents unauthorised users from eavesdropping on the network traffic.

Audit and report cost

This audit and report cost $485,000 to produce.

Copyright

ISSN 1834-1128.

qao.qld.gov.au/reports-resources/parliament

Suggest a performance audit topic
Contribute to a performance audit in progress
Subscribe to news and our blog
Connect with QAO on LinkedIn

T: (07) 3149 6000
M: qao@qao.qld.gov.au
W: qao.qld.gov.au
53 Albert Street, Brisbane Qld 4000
PO Box 15396, City East Qld 4002

Queensland Audit Office
Better public services