ITEM NUMBER: 9.2

CONFIDENTIAL REPORT

IT DISASTER RECOVERY PLAN (Report No: 421/20)

Pursuant to Section 90(2) of the Local Government Act 1999 the Report attached to this agenda and the accompanying documentation is delivered to the Council Members upon the basis that the Council consider the Report and the documents in confidence under Part 3 of the Act, specifically on the basis that Council will receive, discuss or consider:

e. **matters affecting the security of the council, members or employees of the council, or council property, or the safety of any person.**

Recommendation – Exclusion of the Public – Section 90(3)(e) Order

1 That pursuant to Section 90(2) of the *Local Government Act 1999* Council hereby orders that the public be excluded from attendance at this meeting with the exception of the Chief Executive Officer and Staff in attendance at the meeting in order to consider Report No: 421/20 IT Disaster Recovery Plan in confidence.

2. That in accordance with Section 90(3) of the *Local Government Act 1999* Council is satisfied that it is necessary that the public be excluded to consider the information contained in Report No: 421/20 IT Disaster Recovery Plan on the following grounds:

e. pursuant to Section 90(3)(e) of the Act, the information to be received, discussed or considered in relation to this Agenda Item is related to matters affecting the security of the Council.

3. The Council is satisfied, the principle that the meeting be conducted in a place open to the public, has been outweighed by the need to keep the information or discussion confidential.

Item No: **9.2**

Subject: **IT DISASTER RECOVERY PLAN**

Date: 16 December 2020

Written By: General Manager, Strategy and Business Services

General Manager: Strategy and Business Services, Ms P Jackson

SUMMARY

An action from the Control Track Assessment in 2019 was for an IT Disaster Recovery Plan to be reviewed and brought back to the Audit Committee. An agreed action from the Cyber Security Review conducted in September 2020 was to update the IT Disaster Recovery Plan for Council. The IT Disaster Recovery Plan is attached for consideration by the Committee.

RECOMMENDATION

That the Audit Committee:

1. advises Council it has received and considered the IT Disaster Recovery Plan for Council; and

<u>RETAIN IN CONFIDENCE - Section 91(7) Order</u>

2. having considered Agenda Item 9.2 IT Disaster Recovery Plan in confidence under Section 90(2) and (3)(e) of the *Local Government Act 1999*, the Council, pursuant to section 91(7) of that Act orders that the report, attachments and minutes be retained in confidence for a period of 24 months and that the Chief Executive Officer is authorised to release the documents prior to that time if and when all parties to the contract have provided their consent.

COMMUNITY PLAN

A Place that Provides Value for Money

COUNCIL POLICY

Not Applicable.

STATUTORY PROVISIONS

Sec 122 Local Government Act.

BACKGROUND

A disaster recovery plan (DRP) is a documented, structured approach that describes how an organisation can quickly resume work after an unplanned incident. It is an essential part of a business continuity plan. It is applied to the aspects of an organisation that depend on a functioning IT infrastructure. The plan consists of the precautions to minimise the effects of a disaster so the organisation can continue to operate or quickly resume mission-critical functions.

REPORT

An action from the Control Track Assessment in 2019 was for an IT Disaster Recovery Plan to be reviewed and brought back to the Audit Committee. An agreed action from the Cyber Security Review conducted in September 2020 was to update the IT Disaster Recovery Plan for Council.

Administration, with assistance from akto, have updated the Council's IT Disaster Recovery Plan. The plan is included as Attachment 1.

Refer Attachment 1

BUDGET

The costs incurred from the implementation of the agreed actions will be incorporated in existing budget.

LIFE CYCLE COSTS

Life cycle costs associated with the recommendations in the IT Disaster Recovery Plan will be considered as part of annual budget processes.

Attachment 1

IT DISASTER RECOVERY PLAN

OUR PLACE 2030 | CITY OF HOLDFAST BAY

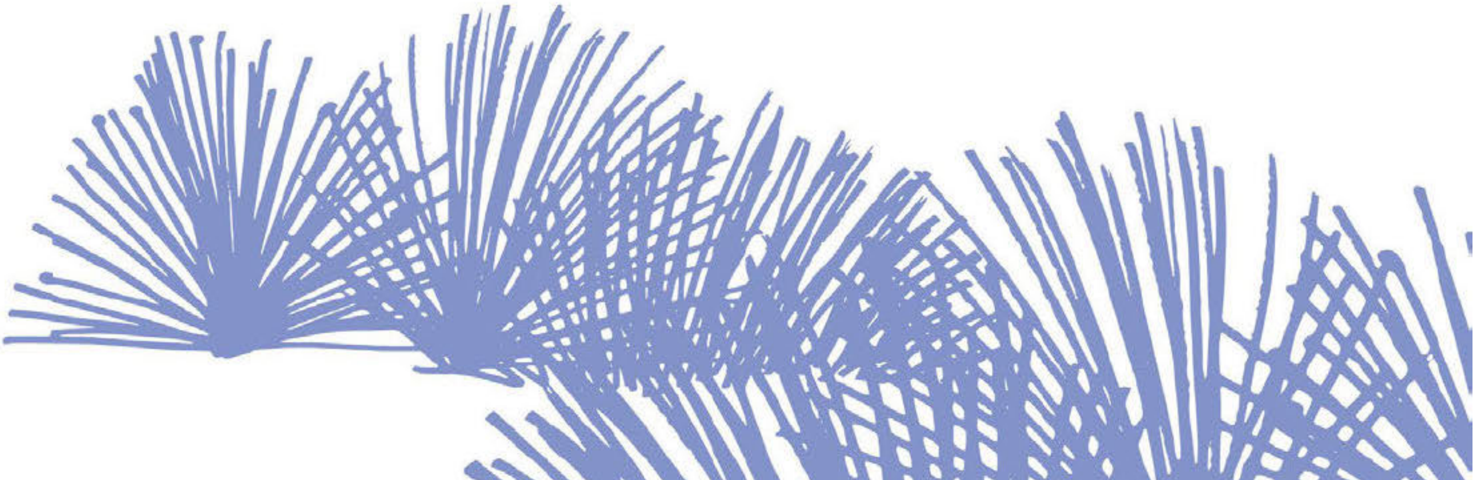

Basket Grass
Lomandra longifolia

TABLE OF CONTENTS

1. INTRODUCTION

An information technology disaster recovery plan (IT DRP) should be developed in conjunction with the City of Holdfast Bay's (CoHB) Business Continuity Management Framework (BCMF). Priorities and recovery time objectives for information technology should be developed during the business impact analysis. Technology recovery strategies should be developed to restore hardware, applications and data in time to meet the needs of the business recovery.

A DR plan should include a complete inventory of hardware and applications in priority order. Each application and piece of hardware should have the vendor technical support contract information and contact numbers, so you can get back up and running quickly. See Table 2 for current Application Matrix.

The starting point of planning is to define the tolerance for downtime and data loss. Figuring out where the business is on the spectrum will determine what type of solution you will need to recover from a disaster. The DRP should divide the applications into six tiers which aligns with the Maximum Acceptable Outage (MAO) states in the Council's BCMF document:

Criticality	Tier 1	Tier 2	Tier 3	Tier 4	Tier 5 not critical	Tier 6 not critical
MAO	Immediate (<4 hrs)	Today (5-24 hrs)	Tomorrow (1-3 days)	This Week (4-7 days)	Next Week (8-14 days)	Eventually (>14 days)

Defining which applications are most important will aid the speed and success of the recovery. The tiers might change based on the results, which could reveal unknown gaps to fill before a true disaster.

CoHB has a comprehensive Council Wide BIA that outlines the following:

- Process criticalities, owners and peak seasons;
- Applications by criticality - Applications list, with filter, RPOs and RTOs;
- Dependencies by criticality; and
- Processes by location, informing loss of site strategy by BU, Dep or Team.

The DRP should be regularly exercised for testing and awareness. It is stated in CoHB's BCMF that all critical functions (Tier 1-4) should be tested every two years and non-critical functions (Tier 5-6) are not required to be tested.

The CoHB have appropriate Crisis Management and Business Continuity Plans that lay out who is responsible for what and identify backup personnel. Having clearly identified roles will garner a universal understanding of what tasks need to be completed and who is responsible for what. This is especially critical when working with third-party vendors or providers. All parties involved need to be aware of each other's responsibilities in order to ensure the DR process operates as efficiently as possible.

Protocols for a DRP must include who and how to contact the appropriate individuals on the DR team, and in what order, to get systems up and running as soon as possible. It is critical to have a list of the DR personnel with the details of their position, responsibilities and emergency contact information. A succession plan should be in place with trained back-up employees in case a key staff member is on holidays / sick leave or in a place where they cannot do their part.

Many times, the main communication platforms (phone and email) may be affected and alternative methods of contacting your employees will be needed. A good communication plan will account for initial communications at the onset of a disaster as well as ongoing updates to keep staff informed throughout the event.

Constituents and customers should be given timely status updates on what they can expect from the business and when. If both your constituents and customers understand that the business is aware of the situation, the organization is adequately prepared and working to take care of it in a timely manner, they will feel much reassured.

CoHB have now migrated many of their key applications to the cloud, a binding agreement with the SaaS providers that defines their level of service in the event of a disaster is essential This will help ensure that they start working on resolving your problem within a specified time. Figure 4 in the Appendix provides an example of this.

DR Plans should be tested reguarly. CoHB's BCMF states that all critical functions (Tier 1-4) should be tested every two years and non-critical functions (Tier 5-6) are not required to be tested. It has been noted that a complete DR test (where Production has been taken offline and DR is being used for production) has not been tested in the last several years.

CoHB need to develop a Test Plan for on-premise backup failovers. It is noted the majority of "core" business applications are now in the cloud, so CoHB should assess effort required versus business risk with respect to on-premise test planning. This can be potentially time, resource and asset intensive to have adequate DR on-premise equipment.

If the DR process is not tested, it is not effective. The backup hardware may have failed, the supply chain may rely on someone incapable of dealing with disaster, the internet connection may be too slow to restore your data in the expected amount of time, the DR key employee may have changed their contact details. There are a lot of things that may break a perfect plan. The only way to find them is to test it when you can afford to fail. The employees that are involved need to be well versed in the plan and be able to perform every task they are assigned to without issue.

Since the last version of CoHB's DRP was released the world has been involved with another type of disaster that the majority DRP's do not cater for – a pandemic.

2. PANDEMICS VERSUS OTHER DISASTERS

Most business continuity planning focuses on events like fires, earthquakes, server crashes, cyber-attacks, and similar disasters that typically cause only a brief disruption until the disaster passes or is resolved, or operations can be relocated to an unaffected area. IT disruptions as well as natural and man-made disasters frequently impact only a particular geographic area, facility or system. By definition, a pandemic (an epidemic or outbreak of infectious diseases that have the ability to spread rapidly over large areas, including worldwide) affects vast numbers of people and, therefore, organisations. As we have seen, the toll on individuals as well as companies and the economy, can be widespread and devastating.

The unpredictable duration of pandemics requires additional considerations, including preparations for continuity of procedures and protocols that might simply be diverted or delayed under other types of disruptions. Other disasters also often only impact a limited number of individuals or processes. COVID-19 has interrupted business at all levels, with few processes unimpacted. Pandemic planning requires everyone to anticipate alternative arrangements and fallbacks for all stages of operations.

Another way in which pandemic planning requires more extensive diligence is that third-party vendors (such as SaaS vendors), and suppliers will also be impacted. It is important that businesses adequately assess the preparedness and resiliency of third-party vendors upon which they rely. A supplier's failure can be incredibly disruptive to business and should be anticipated with mitigation efforts.

What is a pandemic?

A pandemic is a widespread infectious disease that spreads quickly and widely among human or animal populations. Pandemics often relate to a virus such as the H1N1 influenza. Animals frequently develop new viruses and when an animal virus combines with a human one, humans can become ill. Because the virus is new, most of the host population can become highly susceptible to infection.

Why a pandemic is a crisis?

A severe pandemic can disrupt a society and its economy. It could overwhelm a nation's health system and harm its trade. A pandemic may force organisations to take extraordinary staffing measures. The delivery of services and products may have to be adjusted, even stopped. Such actions have financial impacts for businesses and governments alike.

Staffing arrangements during a pandemic may include telecommuting (working remotely from the workplace). This may be necessary to maintain vital services. Equipment, technology and technical support need to be in place and in working order. Staff also need to know how to use facilities such as video conferencing. This may require specialist training.

3. CURRENT STATE

The move to the use of cloud service providers has considerably alleviated the overall business risk of significant failure of the CoHB IT systems. As the CoHB transition more of their processes to cloud service providers more processes will be in place to mitigate the business impact on the CoHB of any potential disaster. However, there are a number which remain on CoHB IT assets, and some limitations on the capability of the CoHB to continue business processing in the event of a disaster.

Several key factors should be noted that have helped in reducing risk and potential loss should a disaster occur:

- Extensive use of Applications as a Service including:
  - o all core council functions currently operated via TechnologyOne's application suite of Property and Rates, Assets, HRP, ECM and soon spatial

    (Intramaps); and
    - Office 365 and Exchange online.
- Use of two sites holding production IT systems in one and backup of all IT systems, together with disk storage and limited processing capability in the second location used for processing of systems not as yet moved "to the cloud";
- The distributed nature of the council business - there are several facilities which could be used as offices, each with reasonable network connectivity to the production and secondary sites;
- The adoption of "Virtual" platform [redacted] and use of [redacted] with similar capacity at the production and DR sites, allows daily copying of all council data and systems from production to the secondary site for the systems not "in the cloud";
- The upgrade to the Citrix environment has allowed for a significant number of concurrent users to access CoHB's applications remotely (ie. from another Council site or if required, from home). This has been demonstrated successfully since the COVID-19 pandemic; and
- [redacted]

With the use of core applications such as TechnologyOne Enterprise Software as a Service and Office 365 in the cloud, the dependence on Council hardware has been reduced, however, network resilience, speed and redundancy has become even more critical to the CoHB's business continuity.

Other than a force majeure event, access to TechnologyOne product is via a hosted Active/Active/Active cloud environment across multiple geographically separated Amazon Web Services instances. The requirement to rebuild and test TechnologyOne product on premise is not required as the disaster avoidance measures in place for high availability of 99.5% uptime is provided as part of the Cloud services agreement with TechnologyOne.

However, it is important that CoHB adequately assess the preparedness and resiliency of all its third-party vendors upon which they rely. COVID-19 has interrupted business at all levels, with few processes unimpacted. Pandemic planning requires everyone to anticipate alternative arrangements and fallbacks for all stages of operation.

The core systems of CoHB are now in the cloud, with the service provider responsible for all backup and DR mechanisms. As a fallback position, CoHB obtains a copy of the data which underpins the business systems, so that in the unlikely event of total failure of TechnologyOne as a supplier, there is a copy of the data available to reconstruct records. It must be understood that:

- the data is not in a form which would allow a rapid return to operations; and
- CoHB would need access to adequate IT processing power to be able to use the data.

The design of IT systems reflects a balance of investment in technology that mitigates risk but does not eliminate all risk. Further steps could be taken to improve network resiliency and onsite power supply. As [redacted] moving the current production site from [redacted] facilities should be investigated.

4. EXISTING DR READINESS

[redacted]

Since late 2015 applications have been progressively moved to the cloud using service providers such as TechnologyOne and Microsoft. As a consequence, the heavy reliance on the critical IT infrastructure is reducing, with networking and authentication measures becoming the key consideration of the maintenance of business continuity.

The following table shows those applications "in the cloud" as Software as a Service (SaaS) vs those applications which remain "On Premises". On premises applications have reduced over the past years, but still manage data or systems of value to CoHB.

[redacted]

For the SaaS systems, redundancy (resilience) of operation is managed by the service provider under the contract, which extends to full Disaster redundancy.

For on premise systems, the existing infrastructure consists of both high availability features [redacted] and several free-standing servers running Citrix at the Production site. The Secondary site, which is the recipient of daily backups, is comprised of a single server with high availability storage. [redacted]

[redacted]

[redacted]

Production systems are backed up to tape, using [redacted] at the Production site, and data is replicated to the DR site at the [redacted] [redacted] Recovery Point Objective and Recovery Time Objectives are daily - that is, if a disaster occurred late in a particular day, all backup (recovered) data could be up to 24 hours out of date, unless systems are lost during backup, in which case the RPO could be 24 hours plus time from scheduled backup to the failure.

[redacted] and there is limited "data anywhere" functionality implemented in the Technology One application suite currently used by Council.

5. COMPUTER SERVERS

Business applications [redacted] or in the Cloud as Technology One Software as a Service I Office 365 I Exchange online. This means that the in-house applications can be run on any [redacted] rather than having to be run on specific server hardware, and core business functions run in the Technology One cloud or Microsoft Cloud are protected by the supplier.

6. SHARED STORAGE

[redacted]

The secondary site is the backup resource for disaster events.

The storage used for terminal services, shared directories, user home directories and application data such as for rates and asset management systems implements two levels of data protection.

- At the production site a series of replicas are taken of the data and associated operating system using [redacted]
- On a daily basis the critical, rapidly changing data is replicated to the secondary [redacted] For less frequently updated files (such as operating systems images) this process is in place on a weekly basis.
- [redacted]
- Core business applications migrated to TechnologyOne Cloud are managed by TechnologyOne, who provide full backup and DRP facility as part of the contracted service.
- CoHB has transitioned to Exchange Online and Office 365 the office productivity files, and email store are moved to the Cloud.

The pre-emptive precautions recommended in the CoHB October 2017 Version 2.1 IT Disaster Recovery Plan are still relevant until the upgrades in the hardware and network infrastructures have been implemented. Once implemented it is highly recommended that multiple targeted Disaster Recovery tests are actioned (ie core applications and network). The learnings of these then be the input in to an updated DRP.

Please see appendix A: Pre-emptive precautions recommended from October 2017 and Current Disaster Recovery Scenarios (Nov 2020).

7. INFRASTRUCTURE REPLACEMENT

Current ICT infrastructure, which includes the following is end of life items:

- Data Centre equipment located at [redacted] and
- Network equipment located at all sites [redacted]

Even though there are up to date support and maintenance contracts, the equipment's age will make it difficult to source replacement parts if there is failure. Currently CoHB are in the process of replacing all network switches. On Premise servers are now 10+ years old and need replacing ASAP. It has been noted that in further investment in moving to the cloud has not been approved and currently CoHB is seeking costs to replace on premise server hardware.

8. CURRENT NETWORK

[redacted] There is no Internet redundancy at Brighton or any other location. Note that while the [redacted]

Telephony at all sites is provided by a [redacted] with a failover system hosted at the [redacted] In-dial telephone number ranges are provided by Telstra SIP, access via the Primary Telstra connection at the [redacted] There is a secondary SIP connection available from a dedicated Telstra SIP service at the [redacted] Server Infrastructure at [redacted] is backed up via [redacted]

<u>There is no redundant link between the two server infrastructure sites.</u>

[redacted]

In January 2020 Data#3 were approached to conduct an on-site review of CoHB's infrastructure to obtain an independent technical assessment. The assessment in scope was to assess CoHB's core Routing, Switching and Wireless infrastructure and form recommendations based on business initiatives such as Cloud adoption of IT services and Smart-Cities.

[redacted]

9. MOVING FORWARD

With the use of core applications such as TechnologyOne Enterprise Software as a Service and Office 365 in the cloud, network resilience, speed and redundancy has become even more critical to the CoHB's business continuity.

With this in mind the following approach is highly recommended.

Phase 1

This phase maintains the current chain topology but increase resilience by ensuring network traffic can flow in either direction if there is link failure.

The following details the high-level activities to be undertaken:

- Network
  - Update current WAN Links from [redacted] sites and introduce redundant links to these site;
  - Establish dual (primary and redundant) paths to Cloud infrastructure to support additional data centre infrastructures when approved, to move to the cloud;
  - SD WAN routers with 4G/5G capable backup at all sites with 4G/5G being enabled at the [redacted]
  - Replace end of life network switches.
- Data centres (Primary and DR)
  - Discovery process to determine costs to house [redacted] data centre in the cloud;
  - Move all [redacted] environments to Cloud infrastructure;
  - Develop and implement Backup, DR and BC strategies based on the new topology.
- Voice
  - SIP
  - [redacted] PABX (both primary and DR)

[redacted]

Phase 2

This phase will continue to utilise the WAN redundancy implemented in phase 1 but will replace the current daisy chain WAN connectivity with a traditional hub and spoke topology. The following details the high-level activities to be undertaken.

- Network (Replace current topology with Hub and Spoke topology)
  - Update WAN Links to enable primary and redundant links for all sites including [redacted]
  - SD WAN capable routers with 4G/5G backup for [redacted]
  - Replace remaining (if any) end of life network switches.

- Data centres (Primary and DR)
  - Opportunity Develop and implement Backup, DR and BC infrastructure and strategies utilising [redacted] ICT facilities due to [redacted]

- Voice
  - Relocate DR SIP services to [redacted]
  - Relocate DR PABX to [redacted]

[redacted]

APPENDIX A

Pre-emptive precautions recommended

While the situations described in this document are quite unlikely, it is probably worthwhile taking some relatively simple pre-emptive steps.

1. In a disaster that requires restore from tape, a working backup server and tape drive is required. While one could be built from scratch in around four to six hours once materials and licenses are at hand, it would be far simpler to have a known good copy stored on an offline volume at the DR site. Similarly, a backup copy of any servers used by IT to administer the system would be useful.
2. A workstation with [redacted] client and copies of documentation, including the DR Plan will be required, and should always be maintained in good working order.
3. Access to an alternative path to the Internet, such as by iPhone tethering will be required until internet connectivity and web access over the network re-enabled.
4. Paper based documentation should be maintained and accessible including server lists, DR Plan and backup tape inventories and recall procedures, independent of the functioning of the primary site. It is recommended that an up-to-date copy should also be stored securely off-site.

Potential Disasters considered

This document addresses disasters from an Information Technology perspective, including:

1. Malicious cyber-attack such as by a worm, Trojan or virus which can result in data loss or damage in production systems;
2. Malicious or accidental acts by trusted personnel resulting in damage to company data assets, which may be the primary production site, secondary site or both;
3. Destruction of the production environment such as by fire, earthquake or some other major event;
4. Destruction of both production and secondary sites such as by earthquake or malicious attack; and
5. Outage of telecommunication / networking facilities perhaps caused by civil disruption, industrial action or natural events.

This plan does not consider risks outside of IT such as loss of key personnel or loss of hardcopy materials held in a single location, matters which should be part of a broader Business Continuity Plan. This plan is not a council wide Disaster Recovery Plan or Business Continuity Plan, which is a wider scope that includes IT Disaster Recovery Planning as a sub-component.

Risk Analysis Table

The table on the following page shows some of the risks that are typically considered when examining Disaster Readiness. The table shows type of risk on the left, along with pre-mitigation consequence.

Since CoHB has already invested considerable expense and effort to minimise risk, many of the risks are reduced, so that most issues are addressed by the use of the primary site / secondary site data replication strategy adopted over the last four years.

Mitigation steps taken already reduce most risks to a level that would generally be seen as acceptable, EXCEPT for the residual risk that relates to unauthorised erasure of data at production and DR, such as by a hacker or a disgruntled employee or contractor - or by innocent misadventure by a network administrator.

It is recommended that tight security and access control is in place to avoid issues such as shared passwords and controlled contractor user ids and passwords to minimise the residual risk of malicious damage. Innocent misadventure can be minimised by adopting a "two person" rule when modifying systems data at storage level.

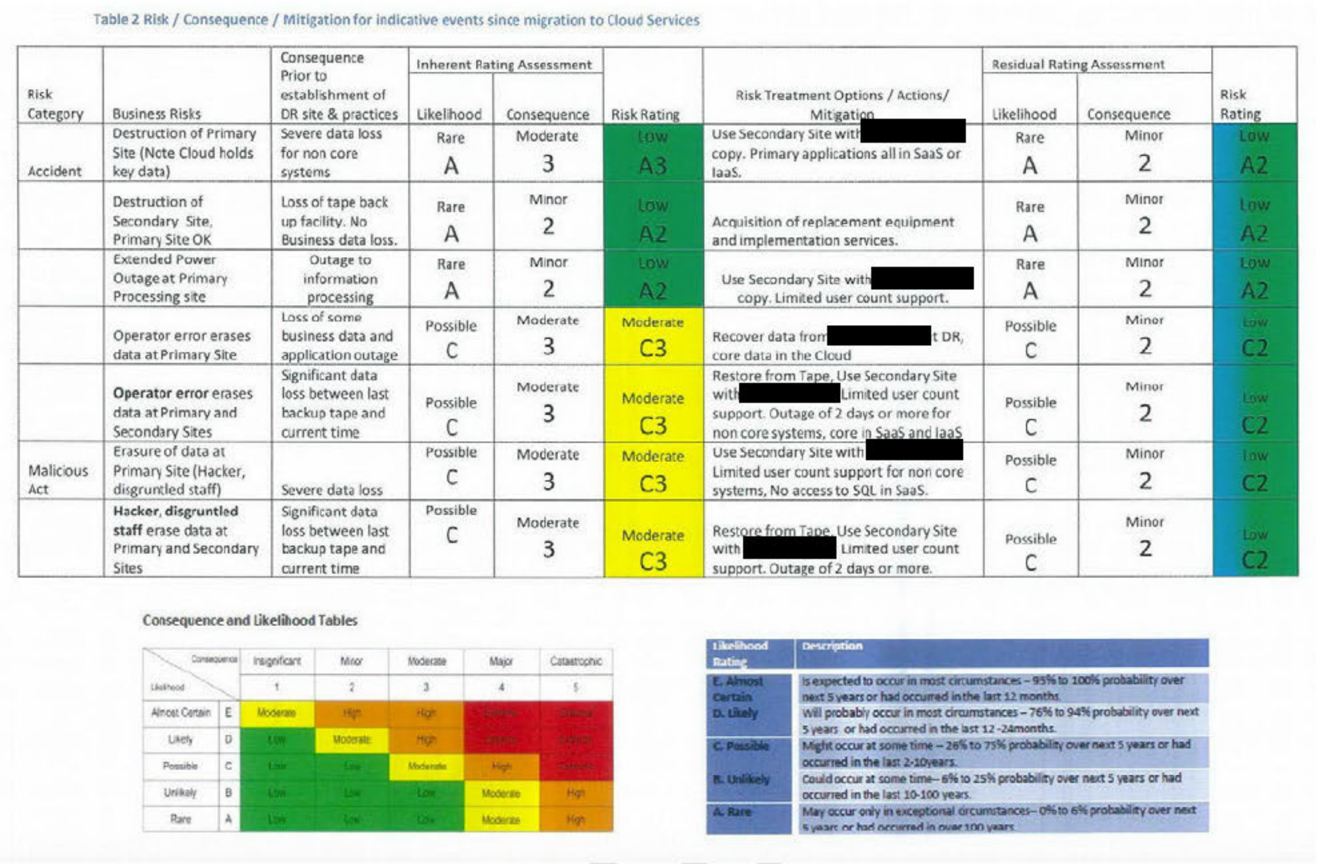

Table 2 Risk / Consequence / Mitigation for indicative events since migration to Cloud Services

Risk Category	Business Risks	Consequence Prior to establishment of DR site & practices	Inherent Rating Assessment		Risk Rating	Risk Treatment Options / Actions/ Mitigation	Residual Rating Assessment		Risk Rating
			Likelihood	Consequence			Likelihood	Consequence	
Accident	Destruction of Primary Site (Note Cloud holds key data)	Severe data loss for non core systems	Rare A	Moderate 3	Low A3	Use Secondary Site with [redacted] copy. Primary applications all in SaaS or IaaS.	Rare A	Minor 2	Low A2
	Destruction of Secondary Site, Primary Site OK	Loss of tape back up facility. No Business data loss.	Rare A	Minor 2	Low A2	Acquisition of replacement equipment and implementation services.	Rare A	Minor 2	Low A2
	Extended Power Outage at Primary Processing site	Outage to information processing	Rare A	Minor 2	Low A2	Use Secondary Site with [redacted] copy. Limited user count support.	Rare A	Minor 2	Low A2
	Operator error erases data at Primary Site	Loss of some business data and application outage	Possible C	Moderate 3	Moderate C3	Recover data from [redacted] t DR, core data in the Cloud	Possible C	Minor 2	Low C2
	Operator error erases data at Primary and Secondary Sites	Significant data loss between last backup tape and current time	Possible C	Moderate 3	Moderate C3	Restore from Tape, Use Secondary Site with [redacted] Limited user count support. Outage of 2 days or more for non core systems, core in SaaS and IaaS	Possible C	Minor 2	Low C2
Malicious Act	Erasure of data at Primary Site (Hacker, disgruntled staff)	Severe data loss	Possible C	Moderate 3	Moderate C3	Use Secondary Site with [redacted] Limited user count support for non core systems, No access to SQL in SaaS.	Possible C	Minor 2	Low C2
	Hacker, disgruntled staff erase data at Primary and Secondary Sites	Significant data loss between last backup tape and current time	Possible C	Moderate 3	Moderate C3	Restore from Tape, Use Secondary Site with [redacted] Limited user count support. Outage of 2 days or more.	Possible C	Minor 2	Low C2

Consequence and Likelihood Tables

Consequence / Likelihood		Insignificant	Minor	Moderate	Major	Catastrophic
		1	2	3	4	5
Almost Certain	E	Moderate	High	High	Extreme	Extreme
Likely	D	Low	Moderate	High	Extreme	Extreme
Possible	C	Low	Low	Moderate	High	Extreme
Unlikely	B	Low	Low	Low	Moderate	High
Rare	A	Low	Low	Low	Moderate	High

Likelihood Rating	Description
E. Almost Certain	Is expected to occur in most circumstances – 95% to 100% probability over next 5 years or had occurred in the last 12 months.
D. Likely	Will probably occur in most circumstances – 76% to 94% probability over next 5 years or had occurred in the last 12 -24months.
C. Possible	Might occur at some time – 26% to 75% probability over next 5 years or had occurred in the last 2-10years.
B. Unlikely	Could occur at some time– 6% to 25% probability over next 5 years or had occurred in the last 10-100 years.
A. Rare	May occur only in exceptional circumstances– 0% to 6% probability over next 5 years or had occurred in over 100 years.

Recovery Point Objectives and Recovery Time Objectives

In discussing Information Technology Disaster Recovery Planning, the term Recovery Point Objective (RPO) is used to refer to the loss of information during the time between the most recent useable backup and a disaster occurring. As there is a trade-off between frequency of systems backup and performance I systems cost, there is generally an (implicit) business decision made as to the data that can be irretrievably lost by the business versus cost of the steps to reduce this time.

Recovery Time Objective is the time between the disaster occurring and the systems being useable again.

It should be noted that information gathered by the organisation and only stored in IT systems during the RPO interval will be missing from the systems which are brought up at the end of the RTO. Experience in other organisations is that much data can be re-assembled from paperwork that would normally have been discarded shortly after being transcribed into the IT system, but any real-time activity may be lost forever during the RPO. Any mission critical data that has a high business impact is a candidate for consideration of techniques such as (non-IT) audio recording for telephony data, or document retention for a short period for any interim paperwork.

Any direct web entry which is critical and unable to be re-assembled should be considered as a candidate for transaction logging to a "Cloud Based" service as a tertiary backup, outside of the general CoHB environment. While these preventive measures have a real cost, they should be assessed against the cost of data loss.

Table 3 Existing Target RPO and RTO

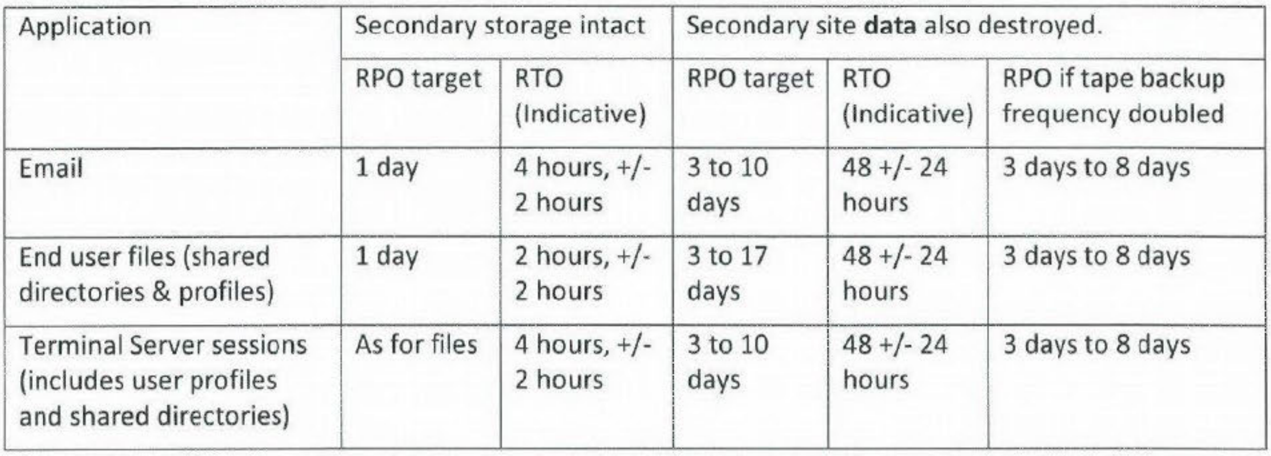

Application	Secondary storage intact		Secondary site **data** also destroyed.		
	RPO target	RTO (Indicative)	RPO target	RTO (Indicative)	RPO if tape backup frequency doubled
Email	1 day	4 hours, +/- 2 hours	3 to 10 days	48 +/- 24 hours	3 days to 8 days
End user files (shared directories & profiles)	1 day	2 hours, +/- 2 hours	3 to 17 days	48 +/- 24 hours	3 days to 8 days
Terminal Server sessions (includes user profiles and shared directories)	As for files	4 hours, +/- 2 hours	3 to 10 days	48 +/- 24 hours	3 days to 8 days

The actual RPO has a range depending on when the last backup was taken and when the disaster occurs. Illustrated in the diagram below for a single site disaster which allows restoration from the Veeam replica of production data. Figure 1shows the disaster happening sometime after completion of a successful Veeam replica cycle. In this case the RPO is less than the "Target RPO."

Figure 2 RPO and "Convenient" Disaster Timing

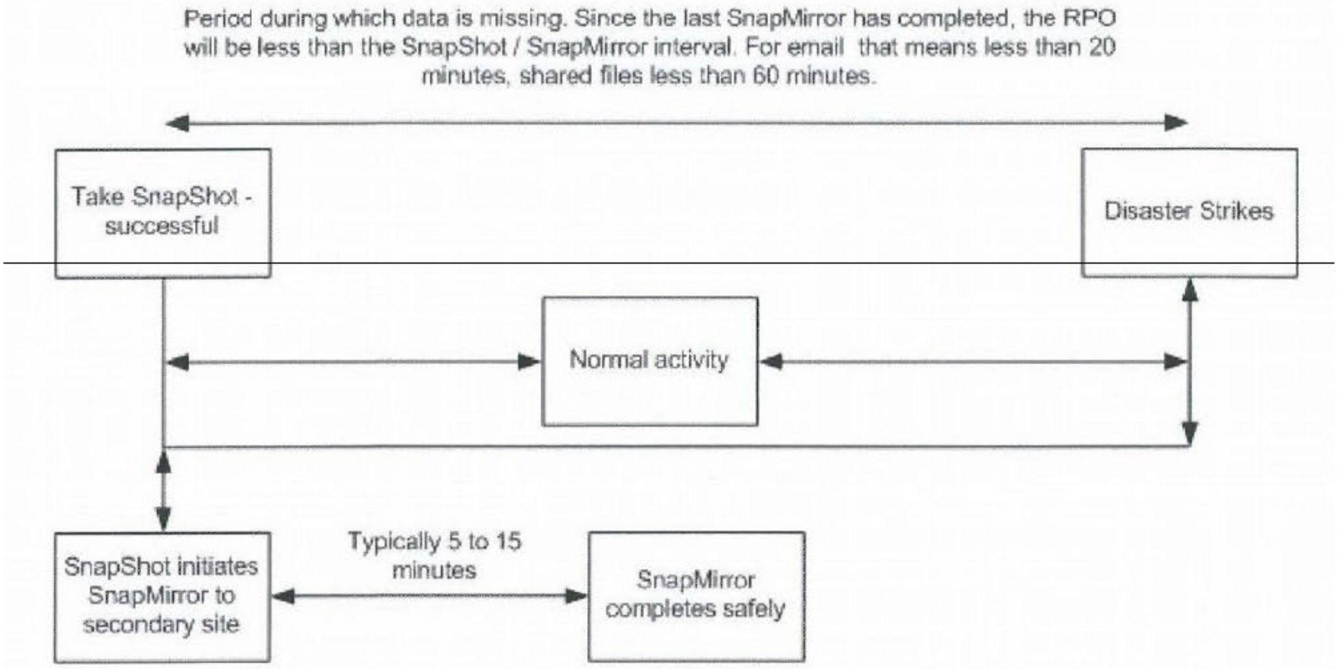

Period during which data is missing. Since the last SnapMirror has completed, the RPO will be less than the SnapShot / SnapMirror interval. For email that means less than 20 minutes, shared files less than 60 minutes.

Take SnapShot - successful
Disaster Strikes
Normal activity
SnapShot initiates SnapMirror to secondary site
Typically 5 to 15 minutes
SnapMirror completes safely

In the above case the "actual RPO" is less than the "target RPO" as the disaster occurs after a Veeam replica cycle, but before the next such cycle.

In the case below, where the disaster occurs before the Veeam replica completes, then the actual RPO would be longer than the target RPO, as the time between the commencement of the Veeam replica process is the "Target RPO."

Figure 3 RPO and “Poor” Disaster Timing

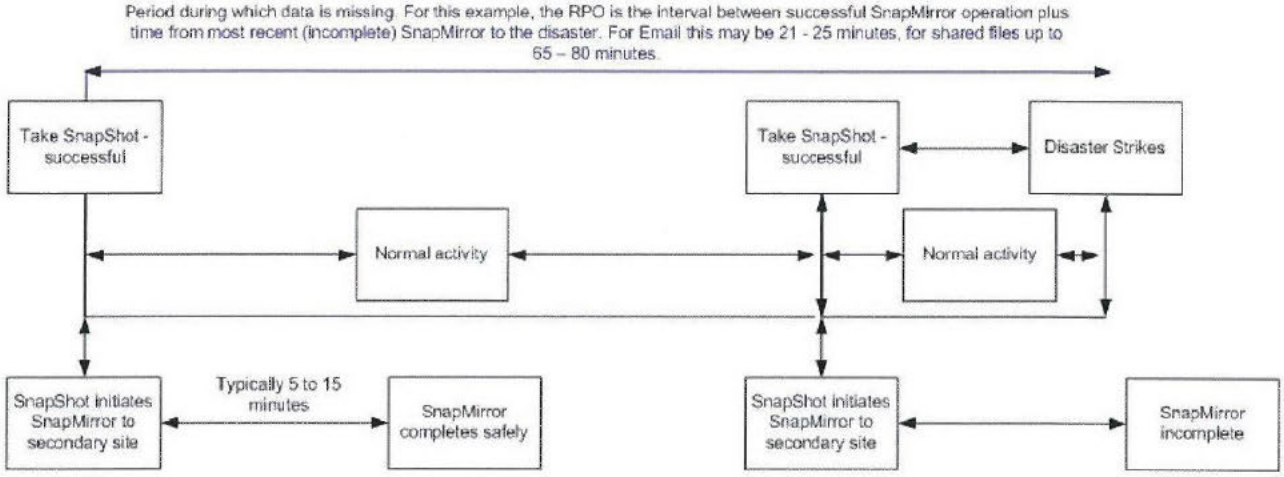

Period during which data is missing. For this example, the RPO is the interval between successful SnapMirror operation plus time from most recent (incomplete) SnapMirror to the disaster. For Email this may be 21 - 25 minutes, for shared files up to 65 – 80 minutes.

Take SnapShot - successful
Normal activity
Take SnapShot - successful
Disaster Strikes
Normal activity
SnapShot initiates SnapMirror to secondary site
Typically 5 to 15 minutes
SnapMirror completes safely
SnapShot initiates SnapMirror to secondary site
SnapMirror incomplete

In this second case, the "actual RPO" will be up to "target RPO" plus Veeam replica time. On current timings that would be one day for e-mail and shared files.

Tape Based Recovery

While the diagrams above are for a Veeam replica recovery, the same is true of tape-based backup, but with a far greater time for data recovery. At the moment the "Backup" phase of the automated weekly backup jobs takes 48 to 72 hours or thereabouts. That is two to three days. That means that the Actual RPO could vary between just over two days to seven days plus just under three days - maybe ten days.

It is possible that optimisation of backup sequence and choice of volumes could reduce the time to backup critical business data. Data which is important to the business could be targeted for frequent backup to tape, while server images which are relatively static could perhaps be backed up less often.

All data that is truly mission critical and requires a low RPO, is now implemented as a "Cloud" solution with full backup and recovery the responsibility of the Could Service Provider.
It should be noted that the business-critical data held in the Tech One Cloud has a RPO and RTO interval as per below.

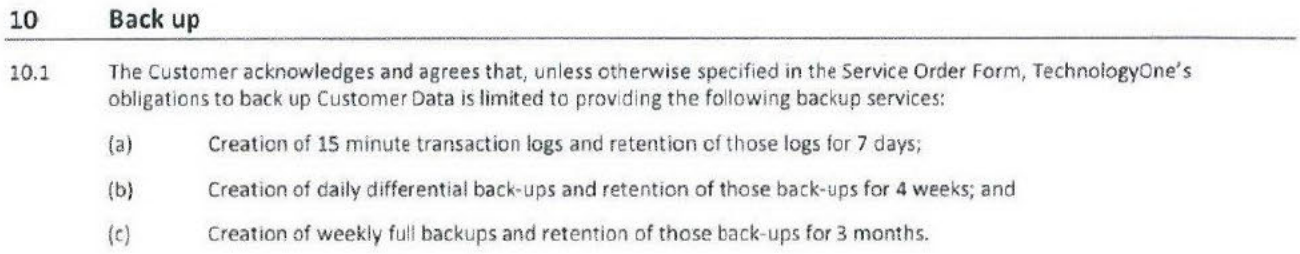

10 **Back up**

10.1 The Customer acknowledges and agrees that, unless otherwise specified in the Service Order Form, TechnologyOne's obligations to back up Customer Data is limited to providing the following backup services:

(a) Creation of 15 minute transaction logs and retention of those logs for 7 days;

(b) Creation of daily differential back-ups and retention of those back-ups for 4 weeks; and

(c) Creation of weekly full backups and retention of those back-ups for 3 months.

Figure 4

DR Processes

The process to be followed after a Disaster Event differs depending on whether the Secondary site data is intact or otherwise, and whether the Primary Site has been physically damaged.

Note that these processes only apply to on premise systems and applications. Cloud (SaaS) applications are "always on" and accessible via internet to the SaaS provider. Core corporate ERP system data is stored predominantly in SaaS systems. For example, Rating, Customer, Property, Infringement, Animal, Financial, HR/Payroll and Email.

The cases considered are tabulated below

Table 4 "Disaster" Cases Documented

Label	Summary	Production Site	Secondary Site	Network
"1A"	Primary Site Physically Destroyed	Destroyed such as by fire, earthquake or flood. Data and equipment unusable	Intact, with data up to date prior to Primary site disaster	Operating normally
"1B"	Primary Site Data destroyed	Data destroyed such as by virus, worm or malicious act. Equipment is functional	Data intact and equipment is functional	Operating normally
"2"	Primary and Secondary Site Destroyed	Data destroyed such as by virus, worm or malicious act. Equipment is functional	Data destroyed such as by virus, worm or malicious act. Equipment is functional	Operating normally and access to Secondary and "site" routine
"3"	Network inoperable	Functioning normally	Functioning normally	Inoperable due to industrial action or geographically wide-spread damage
"4"	Secondary Site Destroyed	Functioning normally	Site destroyed such as by flood. Back-up to tape capability lost	Operating normally and access to Secondary and "site" routine

At least one workstation has been maintained securely at the secondary site which has the following tools at a minimum, to allow initial steps to be taken:

- Network access to the internal network range;
- Web browser, allowing access to web admin for the [redacted]
- An electronic copy of IT information including this document;
- [redacted]

The tools listed above are kept current, so that they will interact with the relevant IT facilities.

Case 1 A - Primary Site Physically Destroyed

While the case 1A and 1B appear similar, they require significantly different actions. If the Primary site including equipment is destroyed, the issues to be faced would be:

- Delay in resuming business services while processing is restored to operation at the secondary site;
- Potential loss of data since the last synchronisation time - with RPO for each application close to the target RPO shown on page 10;
- Low processing power of the secondary site limiting the number of concurrent users supported and IT workload throughput until the Primary Site infrastructure or equivalent processing capacity at the DR site is re-built;
- Re-equipment and integration of a new primary processing site, which is likely to take some days or weeks to accomplish.

To bring the secondary site into use as the primary processing facility, the process to be followed is as outlined below.

Network:

Routing tables (and VPN concentrators if in use) will need to be reconfigured to deliver the address range of production systems into the secondary site, rather than the primary processing site.

Servers and Storage:

[redacted] will need to be enabled for production use. Bring up Active Directory /Domain controllers prior to bringing up Windows file shares.

RPO and RTO - Case 1A

For this case the RPO should be in line with the values shown at above.

RTO is likely to be of the order of two to eight hours, along the following lines:

Shared files, web access, email	→	2 hours at best, to six hours at worst.
Applications (inhouse)	→	4-8 hours but at very low transaction rates
Applications (outsourced)	→	2-6 hours as per web access

It should be noted that IT resources will need to be rationed, as the Secondary Site storage and the single [redacted] which is installed at the secondary site are inadequate to support the normal full workload. It is probable that processing capacity will be less than 50% of normal, and that memory constraints may prevent some application sub-systems from loading and executing.

Case 1B - Primary Site Data Destroyed

If the primary site data is destroyed, either totally or partially - but the equipment is fully functional then the simplest solution is to use [redacted] o re-build the data at the Primary site. The total data transfer time for a full reverse transfer is likely to be of the order of a few days.

While the data is being re-built, the Secondary site could be used as production in the same way as outlined for Case A, although that would impose a large workload on the IT team and may complicate matters.

If the data loss at production occurs just prior to a weekend it is likely that the most sensible path would be to simply use the [redacted] backups to restore the data back to the Primary site, without the added complexity of re-enabling production facilities at the Secondary site.

Network:

If the data on the Primary Site is restored using [redacted] and then servers brought up, no network changes should be required. If the Secondary Site is used as production while the data is being restored to the Primary Site, then the network changes referred to in case 1A will need to be implemented.

Servers and Storage:

Servers should be brought back up along the lines outlined in case 1A, either on the Primary Site or Secondary, as decided at the time. Ensure that a DNS server is accessible to the Primary site prior to attempting to bring up any Windows file shares at the Primary site.

RPO and RTO - Case 1B

For this case the RPO should be in line with the values shown previously.

RTO is likely to be of the order of two to eight hours, along the following lines if the secondary site is used as production while the Primary site data is being rebuilt.

Shared files, web access, email	→	2 hours at best, to six hours at worst.
Applications	→	4-8 hours

It should be noted that IT resources will need to be rationed, as the Secondary Site storage and [redacted] will not support the normal full workload.

If the Data is replicated back to the Primary site before resuming production, the RTO will increase by up to two days. RPO will not change.

Case 2 - Primary and Secondary Site Data Destroyed

If the CoHB data infrastructure suffers a data loss across both Primary and Secondary Sites, but the equipment was operative, the recovery scenario includes the use of backup tapes and replication from the Primary site to the Secondary site to put back in place normal production / DR data integrity measures. The recovery from tape will increase RPO to two or more days, and RTO to more than two days.

If the Primary and Secondary site data is destroyed, the issues to be faced would be:

Recovery of IT processing capability, adequate to support business processes - expected to take two to three few days. Loss of data since the last tape backup - RPO for all applications of at least 1day and up to 2 days.

The overall recovery process involves bringing data back into the Primary site from tape, and then replicating that to the Secondary site using Veeam.
To re-enable the council processing facilities, the process to be followed is as outlined below. Steps and timeframe in this recovery process are:

Obtain backup tape from [redacted]	→	4 hours
Enabled [redacted]	→	2 hours
Mount tape and restore data [redacted]	→	24 - 48 hours

RPO and RTO - Case 2

If data on the primary and secondary sites is lost, then the RPO will be at least one day, and may be as much as two to three days. RTO is likely to be of the order of two or more days.

Case 3 Network Inoperable

The CoHB network has already been modified over a number of years to include a degree of redundancy. Measures already taken by CoHB include the provisioning of redundant paths between Primary Processing Site and Secondary Site, and all remote sites and the Primary and Secondary sites.

The low probability of a severe network outage in any one year, and the difficulty of predicting what circumstance would be faced in such an outage, makes it hard to document any specific actions to be taken, should such an outage occur.

Steps that can be taken include diversifying network carriage service providers and implementing 4G / 5G fall back for remote site routers.
Use of tablets and smart phone devices has been implemented by CoHB and the use of these device to access core operational data should be well understood.
With the transition of core enterprise applications and data to Cloud, the sole reliance on internal network connectivity has reduced and the ability to provide 4G/5G backup connectivity to Cloud services is available to be able to continue access to enterprise applications.

Case 4 -Secondary Site (Equipment) Destroyed

The key issue presented by destruction of the Secondary site equipment is the subsequent exposure to serious data loss if an incident occurs at the Primary site during the time between loss of the Secondary site and acquisition of new Secondary facilities.

This scenario does not require any direct action to continue business application processing. It would however require prompt attention to resume a sound backup regime.

Given the move of core systems to the Cloud Service Providers the low probability of event is such that further documentation is not warranted for this aspect.

Current Disaster Recovery Scenarios (Nov 2020)

Scenario	**Current Mitigation**	**Future 2-3 year plans**
Loss of internet	Failover to [redacted] Secondary mitigation Failover to [redacted] 4G	Installation of NBN at Production [redacted] Installation of [redacted] Installation of [redacted] Re-architecture of network Outlined in Phase 1 and 2 steps above
Loss of production DC	Failover essential networking services to [redacted] Turn on Domain Controllers [redacted] Re-configuration and activation of DHCP Scopes and Cisco IP Helper addresses if required. Utilise TechOne / Office 365 SaaS	
Cyber-attack – Data corruption (eg. Cryptolocker)	Data restoration of services [redacted] Restoration [redacted] [redacted] Restoration from Tape Recall tapes from [redacted] for monthly or yearly archives	

Please refer to – "*Server Startup and Test Process.xlsx"* as of November 2020, attached to this document.