SERVICE DESCRIPTION: READY MESSENGER HIGH AVAILABILITY

PART A - PRODUCT OVERVIEW

Ready Messenger High Availability is an optional premium service available to customers of One New Zealand Ready Messenger that deploy a separate platform instance during maintenance windows.

If there is a conflict between Service Descriptions for other One New Zealand services, then this service description will take precedence.

PART B – READY MESSENGER HIGH AVAILABILITY SERVICE

1. YOUR READY MESSENGER HIGH AVAILABILITY SERVICE

1. 1 The Ready Messenger High Availability service is deployed as an optional chargeable service and available only for customers that can use the Ready Messenger HTTP/S API.

1. 2 Your Ready Messenger High Availability service will be provided access to a separate instance of the platform. The second instance will be leveraged when a maintenance window becomes active or during an incident resulting in an outage on the primary instance. Only the customers of the Ready Messenger High Availability service will have access to the second instance.

1. 3 Ready Messenger High Availability includes geographically diverse (hosted at Rackspace datacentre in Erskine Park and Alexandria NSW Australia - dual sites) failover for the Ready Messenger APIs.

1. 4 The Ready Messenger service will determine platform unavailability (including hardware, network and software incidents), and when detected, the Ready Messenger High Availability service customer traffic will be re-routed to the second instance.

1. 5 As customers using the Ready Messenger High Availability service remain on the standard instance unless it is not available, the independent second instance remains available to be used for any emergency or scheduled outages. Scheduled maintenance on the second instance will only be performed outside any maintenance work scheduled on the first instance and outside of peak business hours.

1. 6 A new domain is to be used by customers enabled on the Ready Messenger High Availability solution for HTTPS API services.

1. 7 During Ready Messenger scheduled maintenance windows this new domain will be manually pointed to the second instance hosted and operated by third party vendor, Soprano Design with no service disruption.

1. 8 During Ready Messenger instances of unscheduled service outages, this new domain will automatically point to the second instance hosted and operated by third party vendor, Soprano Design with up to 2 minutes to process failover.

2. READY MESSENGER ARCHITECTURE DESIGN

2. 1 The second instance that Ready Messenger High Availability customers will use during a scheduled outage is transparent and identical to the standard Ready Messenger production instance.

2. 2 The second instance is installed in a geo-diverse, tier-one hosting facility.

2. 3 Any account provisioned on the second instance uses the same username and password as used for the same account on the primary Ready Messenger instance.

2. 4 The data from the first instance is not replicated to the second instance nor is the data on the second instance replicated to the first instance.

2. 5 Only one-way Mobile Terminated (MT) based services can be enabled on the Ready Messenger High Availability solution.

3. READY MESSENGER HIGH AVAILABILITY SERVICE LIMITATIONS

3. 1 The method of failover implemented currently on the infrastructure that delivers the Ready Messenger service is Automatic Failover (High Availability). The change during a failover process may take up to 2 minutes to propagate to any application connecting to the Ready Messenger platform, depending on that application configuration and caching settings. Customer applications with longer caching settings will experience a long time to failover. It is recommended that the customer does not overwrite the DNS TTL or not exceed 2 minutes for caching DNS lookup.

3. 2 Both automatic and manual methods are implemented on the production infrastructure and are used when required.

3. 3 The Ready Messenger High Availability service will deliver a service availability of 99.99% based upon a rolling 12-month measurement interval.

3. 4 Any applications with longer caching settings will experience a long time to failover.

3. 5 The Ready Messenger High Availability service may use a secondary short code for the period that the second instance is leveraged.

3. 6 The independent second instance remains available to be used for any emergency or scheduled outages, and will only become available for use if the standard is not available.

3. 7 In the event that Disaster Recovery failover is required, a manual failover process will be followed. During such time the platform availability will be impacted and customers will not be able to access the service. To protect you from such scenarios and ensure that you continue to access the service while Disaster Recovery failover is being performed on the platform primary instance, the Ready Messenger High Availability service will protect your service, automatically switching to the second independent instance until service is restored on the primary instance.

3. 8 Ready Messenger monitoring capabilities are designed to auto detect platform unavailability (including hardware, network and software incidents) and when detected the Ready Messenger High Availability service customer traffic will be automatically re-routed to the second instance in under 2 minutes. The first and second instance will not have concurrent maintenance windows.

4. FAILOVER SCENARIOS

#	Scenario	Description	Automated Failover	Time to Process Failover
1	Database Crash	DB failure preventing the application from being able to process messages requiring manual failover to DR site.	Failover to independent HA instance	Up to 2 min
2	Host / Hardware Crash	Hardware failure results in failure to process messages requiring manual failover to DR site.	Failover to independent HA instance	Up to 2 min
3	Network access issue	Hardware and software are healthy but external customer access to the application is prevented by network failure.	Failover to DR on Ready Messenger. If DR is also not available, then failover to independent HA instance	Up to 2 min
4	Scheduled Maintenance	This is a controlled failover ensuring the domain is failed over prior to start the maintenance work.	Failover to HA stack before and during maintenance work and restored back to primary stack once maintenance is completed and tested.	Immediate

#	Scenario	Description	Automated Failover	Time to Process Failover
5	Queuing or downstream provider issue	This is a controlled failover ensuring the domain is failed over when the issue is detected.	Failover to HA stack until queuing depleted or downstream provider issue resolved.	Up to 5 min upon detection

5. TECHNICAL SPECIFICATIONS

5. 1 The DNS A Record has a TTL of 120 seconds. If your application is configured to query DNS using your own caching server, please ensure you do not set a longer TTL as this will increase your exposure to an extended outage and interfere with manual failover during planned maintenance.

5. 2 The TLS (SSL) server is not configured for Mutual Authentication. However, your application must accept our public certificate (and trust chain) for the domain. You may download that certificate and chain through your browser and apply it to your systems as appropriate. This certificate will expire every 2 years and if you have no mechanism to automatically update, then you should put in place procedures to manually check and update.

5. 3 We will provision the HTTP API user account onto the High Availability solution. The customer is to direct their test system to the High Availability URL and verify that traffic can be sent. Once successfully completed, the customer is to then co-ordinate with us to initiate fail over to the High Availability stack and ensure that HTTP messages can still be processed. Once this test is satisfactorily competed, we will revert Ready Messenger High Availability to the primary instance (Ready Messenger). At that point customer can configure production systems to use Ready Messenger High Availability URL.

6. READY MESSENGER HIGH AVAILABILITY SLAS

PRIORITY	DEFINITION	RESPONSE TIME	RESOLUTION TIME	PROGRESS REPORTS
P1	A catastrophic fault, where a minimum of 50% of users at a particular site cannot use a particular Service.	15 minutes	4 hours during Service Hours (or 8 hours during Service Hours where fault is assigned to a third party network operator for resolution).	At 1 hour intervals during Service Hours following our initial response.
P2	Up to 50% of users at a particular site cannot use a particular Service.	30 minutes	8 hours during Service Hours (or 16 hours during Service Hours where fault is assigned to a third party network operator for resolution).	At 2 hourly intervals during Service Hours following our initial response.
P3	Repeated quality problems. Intermittent faults or degraded service. No major service impact.	2 hours	5 Business Days	At 8 hourly intervals during Service Hours following our initial response.

In the table above, Resolution Time means the period elapsed (during Service Hours) from the time the Fault is lodged with One New Zealand to the time that the affected Service is restored to normal operation and Service Hours means 7.00am to 5.30pm Monday to Friday (inclusive), but excluding national public holidays.

PART C – PRICING

Pricing and Invoicing Overview

Refer to the Pricing Schedule for details of the applicable Charges. We will invoice for the Ready Messenger High Availability service on a monthly basis once installation is complete. This may be independent of invoicing for other One New Zealand services.

D – DEFINITIONS

In this Service Description:

Failover is the ability to switch to a redundant or standby computer server, system, hardware component or network upon the failure or abnormal termination of the previously active application, server, system, hardware component, or network.

Instance is a single copy of a running application.