ZYNSTRA TECHNICAL BRIEFING NOTE

Service Level Agreement

CLOUD MANAGED SERVER AVAILABILITY

What is the minimum level of 'availability' that can be expected of a Cloud Managed Server?

Standard (single node) *Availability measured over a calendar quarter*	**99.5%**
High Availability (two node cluster) *Availability measured over a calendar quarter*	**99.9%**

How is the severity and impact (and therefore priority) of an issue represented?

Before explaining what is meant by 'availability', it is important to understand how the priority of an issue is defined based on its severity and impact as the definition of availability relates to certain priorities of issue only.

When a Service Provider raises a request or issue with the Support Desk, a priority is assigned to it based on its perceived urgency which, in the case of issues, is related to its severity and its impact. The definitions used when determining the priority of a request or issue are as follows:

Priority 1 (P1) All of the customer's users are unable to access (a) one or more 'critical software components' on a Cloud Managed Server or (b) one or more IaaS VMs on that server, excluding situations in which a disaster (as defined later in this document) has occurred.

Priority 2 (P2) Any users are unable to access a critical software component or one or more IaaS VMs are not available.

Priority 3 (P3) Any users are experiencing an issue with one or more critical software components but the issue does not materially impact their ability to work; that is, the issue is an inconvenience (and possibly a major inconvenience) but the users can work around the issue for a short period.

Priority 4 (P4) A request for a configuration change to the Cloud Managed Server's software platform that only we can make (because of its complexity or side-effects) or a request for advice/support on usage best practice.

The 'critical software components' on a Cloud Managed Server are those which immediately affect any user's ability to access the services and applications running on the server whose management is the responsibility of Zynstra. They include the Zynstra Active Directory Domain Controller, Network Gateway, Security Gateway and Fileserver but they exclude, for example, the Local Backup Software and Cloud Backup Software, neither of which immediately prevent any user from working effectively. The critical software components also include any Zynstra Managed Applications running on the server as well as any Infrastructure-as-a-Service Virtual Machines on the server but they exclude any third party applications installed by the Service Provider or Customer in those Infrastructure-as-a-Service (IaaS) Virtual Machines as the responsibility for these lies with the Service Provider or Customer.

What do we mean by 'Availability' and 'Service Loss'?

A Cloud Managed Server is deemed to be 'available' if it has not experienced a Priority 1 issue that has not yet been resolved. A Service Loss corresponds to any period of time at any time or the day or night on weekdays or at weekends (that is, 24 x 7) during which a Cloud Managed Server is not 'available' except for periods of time resulting from the following:

- Scheduled patches and upgrades that are periodically applied to all Cloud Managed Servers. These can cause periods of Service Loss outside normal the working day in the time zone in which the server is located. The schedule for these is published in advance to Service Providers and Customers.
- Emergency security patches that very occasionally need to be distributed to Cloud Managed Servers at short notice. The decision to apply these and the timing of their application is at our sole discretion and we will only apply such patches if we believe that either the availability of the server or the integrity or privacy of the Customer's data are at risk.
- Factors beyond our reasonable control including issues caused by (a) the Customer, (b) other technology in the Customer's infrastructure that interacts with the Cloud Managed Server, (c) third parties not contracted to Zynstra such as utility and dependent service providers that fail to provide continuous service (e.g. power, connectivity) or (d) natural disasters and force majeure.

How is the duration of a 'Service Loss' measured?

A Service Loss is deemed to have commenced at the earlier of (a) the Service Provider reporting the issue to our Support Desk and (b) the time at which the Service Loss was detected by our automated monitoring capabilities. A Service Loss ends when all the critical software components have been restored to their correct working state and the Service Provider has been notified of this restoration of service.

SUPPORT SERVICES

What is a Support Service?

A Support Service is a collection of service-based commitments that we make in relation to a Cloud Managed Server. It defines the hours during which support and maintenance services will be available to a Service Provider and the speed with which those services can be expected to be delivered. Two levels of Support Service exist, standard and premium. The standard service is delivered to Service Providers at no additional cost. The premium service attracts an additional fee.

What are the service commitments associated with each Support Service?

The levels of Support Service and the service commitments they comprise are as shown in the table overleaf. Hardware-related commitments in the table apply only if the hardware is being supplied as an integral part of the subscription (that is, it is not being purchased separately) and the hardware commitments are subject to the full terms of the HPE Care Pack upon which the commitments are based. Detailed definitions of the terms used in the table and assumptions and dependencies that relate to them are described after the table.

SUPPORT SERVICE SUMMARY

		Our Commitment
Support Desk Working Hours Hours during which requests or issues can be reported and during which we will endeavour to deal with the request or issue based on its priority	P1	**24 x 7**
	P2, P3, P4	**9 x 5**
Support Desk Response Time The target time within which we will endeavour to start working on fixing or creating a workaround for a reported issue	P1	**1** hour
	P2	**4** working hours
	P3, P4	**1** working day
Software Defect Resolution Time The target period of time within which we will endeavour to work around or fix any issues detected in our Software Platform	P1	**4** working hours
	P2	**1** working day
	P3, P4	**2** working days
Hardware Component Failure Resolution Time The target time within which we will endeavour to fix hardware failures that we have detected or that have been reported to us		**9 x 5** Next Business Day (NBD)
Data Recovery Point Target time period in relation to which data created or changed during that period has been lost as a result of the disaster		**1** working day
Cloud Recovery Time Maximum time to restore the Service in the Cloud on an interim basis following the occurrence of a disaster		**1** working day
On-Premises Server Recovery Time Maximum time to restore the Service on a new server on the customer's premises following the occurrence of a disaster		**10** working days

Support Desk Working Hours

When can Service Providers contact the Support Desk?

The times during which Service Providers can contact the Support Desk by phone, email or via the Support Portal for issues or requests are as specified in the Support Service Summary table above and they depend on the Support Service (standard or premium) that is applicable to the server. The Support Desk can only be contacted by Service Providers, not by their Customers, other than in exceptional circumstances where this has been agreed in advance.

Where the contact times are shown as '9 x 5', this means the hours between 0800 and 1700 from Monday to Friday excluding public holidays in the time zone in which the Cloud Managed Server was purchased and the Service Provider is based.

Where the times are shown as '24 x 7', this means every hour on every day of the week including public holidays in the time zone in which the Cloud Managed Server was purchased and the Service Provider is based.

It should be noted that, in the case of Priority 1 issues, and depending on the Support Service (standard or premium) that is applicable, the Service Provider will need to perform its own assessment of priority before making contact with the Support Desk as only issues that are reasonably believed to be Priority 1 issues can be reported during certain hours of the day.

Support Desk Response Time

How quickly can Service Providers expect the Support Desk to respond?

The Support Desk response time is a measure of the speed with which we will endeavour to (a) understand the nature of any request or issue raised by a Service Provider, (b) allocate a priority level to it based on severity and impact (as described in the priority definitions set out previously) and (c) start to action the request or investigate the issue.

The target response time varies based on the severity and impact of the issue and on the Support Service (standard or premium) that is associated with the Cloud Managed Server in question.

Software Defect Resolution Time

How quickly can Service Providers expect software defects to be resolved?

Software defect resolution is the process of restoring a Cloud Managed Server to its normal operating state following the occurrence of a software defect that affects one or more of its critical software components and renders it or them unusable or inaccessible.

The software defect resolution time is measured from the earlier of (a) the time that the Service Provider reported the fault to us, and (b) the time that we detected the fault using our automated monitoring tools. A software defect is deemed to have been resolved when we have implemented either (a) a workaround to the fault that does not materially hinder user productivity, or (b) a permanent fix to the fault.

Hardware Component Failure Resolution Time

How quickly can Service Providers expect hardware failures to be resolved?

Hardware fault resolution is the process of restoring a Cloud Managed Server to its normal operating state following a hardware component failure. The service commitments associated with the process of hardware fault resolution apply only if the hardware has been supplied as an integral part of the subscription; that is, the hardware was not purchased separately.

Hardware component failures are resolved using the HPE Care Pack that is associated with the Support Service (standard or premium) associated with the Cloud Managed Server. The terms and conditions associated with hardware failure resolution are those published by HPE for its Care Packs.

The hardware component failure resolution time is measured from the earlier of (a) the time that the Service Provider reported the fault to us, and (b) the time that we detected the fault using our automated monitoring tools. It is deemed to have been resolved when the Cloud Managed Server becomes 'available' again.

Data Recovery Point and Recovery Time

What do we mean by a 'disaster'?

A disaster is an event which destroys or damages beyond repair a Cloud Managed Server or renders it inaccessible indefinitely. Examples of disasters include theft, flood or fire.

What do we mean by 'disaster recovery'?

Disaster recovery (or DR) is the process of restoring a Cloud Managed Server to its normal operating state after a disaster with the software and data resident on the server being in its state at the time that the most recently Cloud Backup of that software and data commenced.

What are the key steps in the 'disaster recovery' process?

The first step in the recovery process is the rapid restoration of the Customer's software and data into the Cloud so that it is accessible to the Customer there. Zynstra performs this restoration if requested by the Service Provider and endeavours to complete it within the specified Cloud Recovery Time based on the Support Service that is applicable.

The second step in the recovery process is the full restoration of the Customer's software and data on a new Cloud Managed Server on the Customer's premises (which may differ to the original premises if the premises were also damaged during the disaster event).

- If the Cloud Managed Server hardware was supplied by Zynstra, Zynstra will restore a Cloud Managed Server of identical specification on the Customer's chosen premises, pre-loaded with all the Customer's software and data that were on the Cloud Managed Server at the time the most recently completed Cloud Backup commenced. Zynstra will endeavour to complete this within the specified On-Premise RTO based on the Support Service that is applicable.
- If the Cloud Managed Server hardware was not supplied by Zynstra but was purchased by the Service Provider or Customer, the hardware owner is responsible for the supply of identical new hardware

following the disaster from the same party from whom the original hardware was purchased. Zynstra will restore the Customer’s software and data to the new hardware based on the state of the software and data at the time that the most recently completed Cloud Backup commenced. No commitment to the On-Premises Recovery Time is offered in this case as the speed of restoration will depend on the speed of sourcing of the new server hardware and on the Customer’s available download bandwidth. This is because the restoration can be performed remotely (with much higher available internet download bandwidth) by Zynstra if Zynstra has supplied the hardware - this leads to a much faster restoration time.