Master Services Agreement:

Annexure E: Service Schedule - Hosted Virtual Server Services V10-11

GLOBAL MICRO
Intelligent Technology

This Service Schedule for **Hosted Virtual Server Services V10-11** (the "Service") replaces all previously signed/incorporated version(s) of the Service Schedule for Hosted Hyper-V Services and Hosted Virtual Server Services (if any). It forms part of the Master Services Agreement and Master Services Schedule. Its provisions are an integral part of the Master Services Agreement. Words and expressions defined in the General Conditions and Master Services Schedule shall (unless otherwise defined in this Services Schedule) bear the same meanings where used in this Service Schedule. In this Service Schedule, the following words and phrases shall have the following meanings unless the context otherwise requires:

1. Interpretation

1.1. "**Availability Set**" refers to two or more Virtual Machines deployed across different Fault Domains to avoid a single point of failure.

1.2. "**Azure Compute Unit**" is a concept created by Microsoft to compare compute (CPU) performance across different SKUs.

1.3. "**Fault Domain**" is a collection of servers that share common resources such as power and network connectivity.

1.4. "**IOPS**" means disk performance measured per second in Input / Output Operations.

1.5. **"Maximum Available Minutes"** means the total accumulated minutes during a billing month for all Internet-facing Virtual Machines with two or more instances deployed in the same Availability Set. SP measures the Maximum Available Minutes from when the Customer has started at least two Virtual Machines in the same Availability Set to when the Customer has initiated an action that would stop or delete the Virtual Machines.

1.6. "**SAN**" means a Storage Area Network.

1.7. **"Virtualised Hosting"** means a hosting service on which client servers reside upon hardware infrastructure as virtual machines.

1.8. **"VM**" means a machine that has undergone virtualisation and can be deployed individually or as part of an Availability Set.

1.9. **"VM Profile"** means a specification that defines a VM's memory, CPU, storage capacity, and cost (hourly, monthly or yearly). Some VM Profiles are designed for standard applications, whereas others are for CPU- or memory-intensive applications.

2. Service Overview

2.1. The Service provides access to one or more of the following:

2.1.1. Microsoft Hyper-V virtualisation technology via a hosted cloud-based service;

2.1.2. Showcase Functionality via a hosted cloud-based service.

3. Standard Features

3.1.1. Customers may select from one of the following:

3.1.2. **Unmanaged Virtual Machines** deployed in a single SP data centre. The Customer is responsible for the configuration, patching, and maintenance of the operating system;

3.1.3. **Managed Virtual Machines** provisioned in a single SP data centre. The SP is responsible for the configuration, patching, and maintenance of the operating system;

3.1.4. **Managed Replica Virtual Machines** – VMs deployed and replicated between dual SP South African datacentres where SP is responsible for configuration, patching and maintaining the operating system.

3.2. **Support**

3.2.1. Support for (1) configuration, patching, and maintenance of the operating system of an Unmanaged Virtual Machine, (2) support for any applications installed on any virtual machine, (3) management and configuration of any Anti-Virus and Anti-Malware scanning, and (4) Setup and configuration of Server, is not included as a standard component and is provided via a **Reserved Support Services Subscription** for an additional fee.

3.2.1.1. Customers who wish to log service requests directly with SP described in 3.2.1 above must conclude a separate Reserved Support Services Schedule and associated Reserved Support Services Subscription**.**

4. Performance considerations

4.1. **Virtual Machine Profiles**

4.1.1. A Virtual Machine Profile determines the host computer's hardware by a VM.

4.1.2. Each profile offers different compute and memory capabilities.

4.1.3. The selection of a VM Profile determines the amount of memory and computing power required for the application or software running within the VM.

4.1.4. Virtual Machine Profiles are provisioned on different underlying hardware to cater to performance requirements:

4.2. **Enterprise-grade applications –Dv2 Series**

4.2.1. These profiles are ideal for applications that demand faster CPUs, better local disk performance, or higher memory demands. They offer a powerful combination suitable for many enterprise-grade applications. Dv2-series, a follow-on to the original D-series, features a more powerful CPU.

4.3. **Azure Compute Unit**

4.3.1. The underlying hardware of each series directly impacts the computing performance of VM Profiles.

4.3.2. To compare compute (CPU) performance across SKUs, Microsoft created the Azure Compute Unit (ACU) concept.

4.3.3. ACU can be used to identify which SKU is most likely to satisfy a performance need. ACU is currently standardized on a Small (Standard_A1) VM, being 100, and all other SKUs represent approximately how much faster that SKU can run a standard benchmark.

4.3.4. Note: The ACU is only a guideline. The results for a workload may vary.

SKU Family	ACU/Core – Microsoft Azure VMs	ACU/Core – SPVMs
Standard_A1-7	100	100 – HP ProLiant G7
Standard_A8-11	225*	225* – HP ProLiant G9
D1-14	160	160 – HP ProLiant G8
Dv2 1-14	210 to 250 *	210 to 250* – HP ProLiant G9

ACUs marked with a * use Intel® Turbo technology to increase CPU frequency and provide a performance boost. The Boost amount can vary based on the VM size, workload, and other workloads running on the same host.

5. Premium Storage Disk Limits

5.1. All VMs use persistent Premium Storage. The input/output operations per second (IOPS) and throughput (bandwidth) depend on the disk's size. Currently, there are three Premium Storage disks: P10, P20, and P30. Each one has specific limits for IOPS and throughput, as specified in the following table:

Premium Storage Disk Type	**P10**	**P20**	**P30**
Disk Size	128 GiB	512 GiB	1024 GiB (1 TB)
IOPS per disk	500	2300	5000
Throughput per disk	100 MB per second	150 MB per second	200 MB per second

6. IP Addressing and Internet-Facing Traffic

6.1. SP recommends that all IPs used with the Service are reserved to ensure that (1) an IP address will not change even when stopping or deallocating VMs and (2) that inbound and outbound traffic uses a predictable IP address.

6.2. SP aggregates inbound and outbound traffic to determine internet traffic usage according to the Services Fees Schedule.

6.2.1. SP's traffic measurement system will be the sole method of determining bandwidth and traffic utilisation.

6.2.2. Internet-facing traffic is a best-effort service.

7. Backup and recovery

7.1. Managed Replica Virtual Machines provide high availability. They should not, however, be used instead of a backup.

7.1.1. The Customer must conclude a separate Hosted Backup Services Schedule and associated Hosted Backup Services Subscription for Virtual Machines backup.

8. Managed Virtual Machines

8.1. Managed Virtual Machines and Managed Replica Virtual Machine subscriptions include support for (1) configuration, patching, and maintenance of the operating system and (2) management and configuration of anti-virus and anti-malware scanning.

8.2. SP provides Support during Bronze Coverage hours.

8.3. SP will install monitoring agents on all Windows Virtual Machines.

8.4. SP will deploy and configure anti-virus security on all Virtual Machines.

8.5. SP will respond to the following alerts the monitoring agents raise within Bronze Coverage hours.

Monitor	**Frequency**
Monitor CPU utilization for any unfamiliar usage patterns	every 15 minutes
Monitor memory utilization for unfamiliar usage patterns	every 15 minutes
Monitor hard drive health and space consumption	every 15 minutes
Ensure that all core operating system services are running	every 60 minutes
Monitor Server Agent availability	Every 15 minutes

8.6. Customer shall be entitled to **up to 1000 (one thousand) Pooled Support Units per month per Virtual machine** ('**Pooled Support Allocation**').

8.6.1. SP will only allocate Pooled Support Units after installing monitoring agent(s), and SP has verified that the minimum of the Managed Virtual Machines meets SP's Minimum Requirements.

8.6.2. After that, Pooled Support Units are decremented from the Support Unit Allocation for work performed in response to monitoring agent alerts and may be used on an as-needed basis or according to a customer-requested recurring maintenance plan for the following:

Customer Request	Frequency
Apply Service packs, patches and hotfixes per company policy	Per Customer Request
Reboot servers not responding or inaccessible.	Per Customer Request
Confirm that hypervisor layer antivirus definitions are updated	Per Customer Request
Confirm that virus scans have occurred.	Per Customer Request
Quarantine and Clean viruses from Server(s)	Per Customer Request
Reports of work accomplished and in progress	Per Customer Request
Check event logs for all servers and identify any potential issues	Per Customer Request
Logical Solution Architecture and deployment of Operating System upgrades	Per Customer Request
Optimise disk performance through Disk defragmentation	Per Customer Request
Setup, configuration and design of backups	Per Customer Request
Adhoc assistance with backups and restore activities	Per Customer Request
Formulation of a Backup & Restore Strategy, Documentation of Recovery Point Objectives and Recovery Time Objectives, Scheduled Test restores and disaster recovery testing.	Excluded Requires separate Backup Services subscription
Test or initiate replication failover for Managed Replica Virtual Machines	Per Customer Request
Creation of Customer Specific Monitors or Remediation Scripts	Per Customer Request

8.6.3. Where services are (a) outside the scope of this Service Schedule, (b) relate to Managed Virtual Machine exclusions, (c) are rendered outside of Coverage Hours, or (d) performed by a Work Role other than Tier 2 Support, SP may levy additional fees (based on the Base Labour Rate or BLR specified in the Service Fees Schedule) per hour together with an Uplift per Work Type and Work Roles described below:

Uplift Table	%
Work Types	
Pro-Active Services	+0%
After-hours support Weekdays 18h00-08h00	+25%
After-hours support Weekends & Public Holidays	+50%
Escalation to Microsoft / Citrix / McAfee / VMware / Cisco or another Vendor	+50%
Work Roles	
Skills Development or Trainee	-50%
Tier 1 Support	-25%
Tier 2 Support	+0%
Tier 3 Support	+50%
Tier 1 Developer	+50%
Tier 2 Developer	+100%
Tier 1 Security Analyst	+100%
Tier 2 Security Analyst	+150%
Tier 3 Security Analyst	+200%
Project Manager	+50%
Systems Architect	+150%
Director	+200%

8.6.4. **Measurement Increments -** Labour rendered telephonically or remotely will be measured in increments of 15 minutes and rounded UP to the nearest quarter-hour.

8.6.5. **Onsite Support –** All onsite support is outside the scope of Managed Virtual Machines.

8.6.6. All services requested of the SP that fall outside the terms of these Managed Virtual Machines as described in this schedule are billable.

8.7. **Minimum Requirements for Managed Virtual Machines**

The following is required for SP to provide Managed Virtual Machines:

8.7.1. All Virtual Machines must run a Microsoft Supported version of Windows Server and install all Microsoft Service Packs and Critical Updates (released 30 days earlier).

8.7.2. All server software installed on the Virtual Machines must be genuine, licensed and vendor-supported.

8.7.3. The environment must have a licensed, vendor-supported hardware firewall between the internal network and the internet or an SP-configured firewall.

8.7.4. The Services exclude the fees and costs required to bring Virtual Machines up to these Minimum Standards.

8.7.5. These Minimum Standards are subject to change upon prior written notice from SP.

8.7.6. The Customer agrees to maintain and upgrade the environment to comply with these standards.

Managed Virtual Machines Exclusions

8.8. The Service expressly excludes the following:

8.8.1. The cost of any software, licensing, or software renewal or upgrade fees of any kind;

8.8.2. The cost of any 3rd party vendor or manufacturer support or incident fees of any kind;

8.8.3. The cost to bring Virtual Machines up to Minimum Standards required for Services;

8.8.4. Maintenance of applications software packages, whether acquired from SP or any other source;

8.8.5. Programming (modification of software code) and program (software) maintenance;

8.8.6. Training Services of any kind.

9. Replication across different Fault Domains to avoid a single point of failure

9.1. Managed Replica VMs include replicating a second data centre in a different Fault Domain.

9.2. Replication runs every 15 minutes, assuming the previously scheduled replication has been completed.

9.2.1. Replication depends on the volume of the changes and cannot be guaranteed to be completed within a pre-determined interval.

9.2.2. SP provides bandwidth for the replication up to a maximum of 3% (three per cent) of the disk capacity in any given 12 (twelve) hour period (‘**VM Replication Quota’**). SP may charge extra for any bandwidth over the VM Replication Quota.

9.3. SP will retain up to 24 (twenty-four) replication snapshots.

9.4. **Requirements for Replication:**

9.4.1. Only Managed Virtual Machine support replication.

9.4.2. Virtual Machines must use Geo-Redundant Storage.

10. Service Availability

10.1. If the Service is unavailable, it must be reported to the SP and acknowledged by the SP.

10.2. SP calculates Downtime from when (1) the fault is reported, (2) SP has issued a fault report reference, and (3) SP has acknowledged this as a fault in the Service.

10.3. Following investigation and repair, SP will advise the service restoration time, which will be the end of the Downtime unless the fix is not confirmed.

10.4. **Managed Replica Virtual Machine**

10.4.1. “**Downtime**” means the total accumulated minutes that are part of the Maximum Available Minutes with no External Connectivity.

10.4.2. "**Monthly Uptime Percentage**" is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes - Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Where Downtime is measured in minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of Virtual Machines impacted by that Incident

10.4.3. Exclusions to Service Availability Guarantee:

10.4.4. Any incident which lasts less than 15 (fifteen) minutes;

10.5. **Service Credit:**

Monthly Uptime Percentage	Downtime per month	Bronze SLA Service Credit	Silver SLA Service Credit	Gold SLA Service Credit	Platinum SLA Service Credit
< 99.95%	21.56 minutes	No Credit	No Credit	No Credit	10%
< 99 %	43.2 minutes	No Credit	No Credit	25%	
< 95 %	36 hours	No Credit	100%		