



This Service Schedule for **Hosted Virtual Server Services v8.0.0** (the "Service") replaces all previously signed / incorporated version(s) of the Service Schedule for Hosted Hyper-V Services and/or Hosted Virtual Server Services (if any) and 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**" means a concept created by Microsoft to provide a way of comparing 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 a measure of disk performance measured in Input / Output Operations per seconds.
- 1.5. "**Maximum Available Minutes**" means the total accumulated minutes during a billing month for all Internet facing Virtual Machines that have two or more instances deployed in the same Availability Set. Maximum Available Minutes is measured from when at least two Virtual Machines in the same Availability Set have both been started resultant from action initiated by Customer to the time Customer has initiated an action that would result in stopping or deleting the Virtual Machines.
- 1.6. "**SAN**" means a Storage Area Network.
- 1.7. "**Virtualised Hosting**" means a hosting service in which client servers reside upon a hardware infrastructure as virtual machines.
- 1.8. "**VM**" means a machine than has undergone Virtualisation that can be deployed individually or as part of an Availability Set. .
- 1.9. "**VM Profile**" means a specification that defines the memory, CPU, storage capacity, and cost (hourly, monthly or yearly) for a VM. Some VM Profiles are designed for standard applications, whereas others are designed for CPU-intensive applications, 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. Microsoft Azure Services via a hosted cloud-based service.
 - 2.1.3. Showcase Functionality via a hosted cloud-based service.

3. Standard Features

- 3.1. Virtual Machines can be provisioned in:
 - 3.1.1. **SP's South African datacentres ('ZA VMs')** or
 - 3.1.2. **Microsoft's European datacentres ('EU VMs')**.
- 3.2. **ZA VMs:**
 - 3.2.1. Customers may select from one of the following: (1) **Unmanaged Virtual Machines** – VMs deployed in a single SP South African datacentre where Customer is responsible for configuration, patching, and maintenance of the operating system (2) **Managed Virtual Machines** – VMs protected by 5Nine Cloud Security, deployed in a single SP South African datacentre, where SP is responsible for configuration, patching, and maintenance of the operating system and 5Nine Cloud Security. (3) **Managed Replica Virtual Machines** – VMs protected by 5Nine Cloud Security, deployed and replicated between dual SP South African datacentres where SP is responsible for configuration, patching, and maintenance of the operating system and 5Nine Cloud Security.
- 3.3. **EU VMs:**
 - 3.3.1. Customers may select from one of the following **Microsoft Online Services** which includes the following features: (1) **Managed Virtual Machines** – VMs deployed in a single Microsoft European datacentre where SP is responsible for configuration, patching, and maintenance



of the operating system (3) **Managed Replica Virtual Machines** – VMs deployed and replicated between dual Microsoft European datacentres where SP is responsible for configuration, patching, and maintenance of the operating system.

3.4. Support

3.4.1. 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, and (3) management and configuration of any 5nine Cloud Security's Kernel Mode Firewall, Anti-Virus and Anti-Malware scanning, is not included as a standard component and is provided via a **Support Services Subscription** for an additional fee.

3.4.1.1. Customers who wish to log service requests directly with SP described in 3.4.1 above, are required to conclude a separate Support Services Schedule and associated Support Services Subscription.

4. Performance considerations

4.1. Virtual Machine Profiles

4.1.1. A Virtual Machine Profile essentially determines the hardware of the host computer by a VM.

4.1.2. Each profile offers different compute and memory capabilities.

4.1.3. Selection of a VM Profile is based upon 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 for performance requirements:

4.2. General purpose compute – A1-7 Series

4.2.1. A-series Virtual Machines are deployed on a variety of hardware types and processors. The size is throttled, based upon the hardware, to offer consistent processor performance for the running instance, regardless of the hardware it is deployed on.

4.3. Compute intensive – A8-11 Series

4.3.1. The hardware that runs these Virtual Machines is designed and optimized for compute-intensive and network-intensive applications, including modelling, and simulations.

4.4. Enterprise-grade applications – D and Dv2 Series

4.4.1. These profiles are ideal for applications that demand faster CPUs, better local disk performance, or have higher memory demands. They offer a powerful combination for many enterprise-grade applications. Dv2-series, a follow-on to the original D-series, features a more powerful CPU. The Dv2-series CPU is about 35% (thirty five percent) faster than the D-series CPU. It is based on the latest generation 2.4 GHz Intel Xeon® E5-26xx v3 processors, and with the Intel Turbo Boost Technology 2.0, can go up to 3.2 GHz. The Dv2-series has the same memory and disk configurations as the D-series.

4.5. Azure Compute Unit

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

4.5.2. To provide a way of comparing compute (CPU) performance across SKUs, Microsoft created a concept of the Azure Compute Unit (ACU).

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

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



SKU FAMILY	ACU/CORE – EU VMS	ACU/CORE – ZA VMS
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 amount of the boost can vary based on the VM size, workload, and other workloads running on the same host.

5. Premium Storage Disk Limits

5.1. To provide standardisation between EU and ZA VMs, all disks are provisioned with persistent Premium Storage. The input/output operations per second (IOPS) and throughput (bandwidth) depends on the size of the disk. Currently, there are three types of 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

5.2. **Disk Size:**

5.2.1. All VMs are deployed with a 100GB Primary OS Disk. The remaining profile storage is attached via one or more data disks.

5.2.2. IOPS for EU VMs is reserved by rounding up the Disk size to the nearest Premium Storage Disk option as specified in the table. For example, a disk of size 100 GiB is classified as a P10 option and can perform up to 500 IO units per second, and with up to 100 MB per second throughput. Similarly, a disk of size 400 GiB is classified as a P20 option, and can perform up to 2300 IO units per second and up to 150 MB per second throughput.

5.2.3. ZA VMs are deployed on fibre channel SAN which provided more consistent disk performance than the JBOD storage technologies used in EU VMs. Therefore, all ZA VMs are deployed with P10 Premium Storage Disk Types.

6. Size Tables

6.1. The following tables show the size and capacities various Virtual Machine Profiles provide.

6.1.1. The Storage capacity is represented by using 1024^3 bytes as the unit of measurement for GB.

6.1.2. Maximum network bandwidth is the maximum aggregated bandwidth allocated and assigned per VM type. The maximum bandwidth provides guidance for selecting the right VM type to ensure adequate network capacity is available. When moving between Low, Moderate, High and Very High, the throughput will increase accordingly. Actual network performance will depend on many factors including network and application loads, and application network settings.



6.2. Standard tier: A-Series – general purpose instances

SIZE	CPU CORES	MEMORY	NICS (MAX)	PRIMARY OS DISK	DEFAULT DATA DISK	MAX DISKS	MAX. IOPS PER DISK	MAX NETWORK BANDWIDTH
STANDARD_A1	1	1.75 GB	1	100GB	-	2	2 X 500	MODERATE
STANDARD_A2	2	3.5 GB	1	100GB	35 GB	4	4 X 500	MODERATE
STANDARD_A3	4	7 GB	2	100GB	185 GB	8	8 X 500	HIGH
STANDARD_A4	8	14 GB	4	100GB	505 GB	16	16 X 500	HIGH
STANDARD_A5	2	14 GB	1	100GB	35 GB	4	4 X 500	MODERATE
STANDARD_A6	4	28 GB	2	100GB	185 GB	8	8 X 500	HIGH
STANDARD_A7	8	56 GB	4	100GB	505 GB	16	16 X 500	HIGH

6.3. Standard tier: A-Series – compute intensive instances

SIZE	CPU CORES	MEMORY	NICS (MAX)	PRIMARY OS DISK	DEFAULT DATA DISK	MAX DISKS	MAX. IOPS PER DISK	MAX NETWORK BANDWIDTH
STANDARD_A8	8	56 GB	2	100GB	282 GB	16	8 X 500	HIGH
STANDARD_A9	16	112 GB	4	100GB	282 GB	16	16 X 500	VERY HIGH
STANDARD_A10	8	56 GB	2	100GB	282 GB	16	16 X 500	HIGH
STANDARD_A11	16	112 GB	4	100GB	282 GB	16	16 X 500	VERY HIGH

6.4. Standard tier: Dv2/D-Series – enterprise instances

SIZE	CPU CORES	MEMORY	NICS (MAX)	PRIMARY OS DISK	DEFAULT DATA DISK	MAX DISKS	MAX. IOPS PER DISK	MAX NETWORK BANDWIDTH
STANDARD_DV2/D1	1	3.5 GB	1	100GB	-	2	2 X 500	MODERATE
STANDARD_DV2/D2	2	7 GB	2	100GB	35 GB	4	4 X 500	HIGH
STANDARD_DV2/D3	4	14 GB	4	100GB	185 GB	8	8 X 500	HIGH
STANDARD_DV2/D4	8	28 GB	8	100GB	505 GB	16	16 X 500	HIGH
STANDARD_DV2/D11	2	14 GB	2	100GB	35 GB	4	4 X 500	HIGH
STANDARD_DV2/D12	4	28 GB	4	100GB	185 GB	8	8 X 500	HIGH
STANDARD_DV2/D13	8	56 GB	8	100GB	505 GB	16	16 X 500	HIGH
STANDARD_DV2/D14	16	112 GB	8	100GB	505 GB	32	32 X 500	VERY HIGH

7. IP Addressing and Internet Facing Traffic

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

7.2. Inbound and outbound traffic is aggregated to determine internet traffic usage according to the fees specified in the Services Fees Schedule.

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

7.2.2. Internet facing traffic is a best effort service. If dedicated internet or VPN facing connectivity is required, additional agreements will need to be concluded with the SP.



8. Backup and recovery

- 8.1. Managed Replica Virtual Machines are designed to provide high-availability and should not be relied upon as a backup.
- 8.2. Backup and recovery of Virtual Machines is specifically excluded from this Services Schedule.
- 8.2.1. It is a requirement to conclude a separate Hosted Backup Services Schedule and associated Hosted Backup Services Subscription if backup and recovery of Virtual Machines is a requirement.

9. Showcase Functionality

- 9.1. **5Nine Cloud Security (Only available for ZA VMs)**

An agentless security and compliance solution built specifically for Service Provider Virtual Machines, utilizing the extensibility of the Hyper-V switch which:

 - 9.1.1. **Secures multi-tenant Hyper-V environments and provides VM isolation**
 - 9.1.2. **Protects Hyper-V with fast, agentless antivirus and malware scanning with**
 - 9.1.2.1. Support for Fast incremental scans of Virtual Machines;
 - 9.1.2.2. Support for Orchestrated scans and performance thresholds set across VMs;
 - 9.1.2.3. Support for Staggered scanning for optimal performance;
 - 9.1.2.4. Support for Real-time AV protection.
 - 9.1.3. **Enforces PCI-DSS, HIPAA and Sarbanes-Oxley compliance via a Kernel Model Virtual Firewall with**
 - 9.1.3.1. Support for MAC address and broadcast filtering;
 - 9.1.3.2. Support for ARP rules;
 - 9.1.3.3. Support for Stateful packet inspection;
 - 9.1.3.4. Support for Packet inspection/network traffic anomaly analysis;
 - 9.1.3.5. Support for Inbound and outbound per-VM Quality of Service and bandwidth throttling.
- 9.2. **F5 BIG IP Virtual Edition**

A strong security footprint for applications hosted entirely in SP or Microsoft datacentre.

 - 9.2.1. Dedicated Partition and route domain on a shared Virtual Appliance.
 - 9.2.1.1. Includes: BIG-IP Local Traffic Manager module.
 - 9.2.1.2. Includes: BIG-IP Advanced Firewall Manager.
 - 9.2.1.3. Includes: BIG-IP Application Visibility and Reporting.
 - 9.2.1.4. Includes: SSL Offloading.
 - 9.2.1.5. Support for BIG-IP Access Policy Manager (Subject to additional charges).
 - 9.2.1.6. Support for BIG-IP Application Security Manager (Subject to additional charges).

10. Managed Virtual Machines

- 10.1. Managed Virtual Machines and Managed Replica virtual machines subscriptions include support for (1) **configuration, patching and maintenance of the operating system** as well as (2) **management and configuration of 5Nine Cloud Security's Kernel Mode Firewall, Anti-Virus and Anti-Malware scanning.**
- 10.2. Support will be provided during Silver Coverage hours.
- 10.3. SP will install monitoring agents on all Windows Virtual Machines.
- 10.4. SP will deploy and configure 5Nine Cloud Security on all Virtual Machines.
- 10.5. SP will respond to the following alerts raised by the monitoring agents within Silver Coverage hours.

Master Services Agreement:

Annexure D: Service Schedule - Hosted Virtual Servers v8.0.0



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

10.6. Customer shall be entitled to **up to 3000 (three thousand) Support Units per month per Virtual machine** ('Support Allocation'), where 1000 (one thousand) Support Units is equal to 1 (one) hour of Tier-2 Support during Silver Coverage Hours.

10.6.1. Support Units will only be allocated after the monitoring agent(s) have been installed successfully and SP has verified that the minimum requirements for Managed Virtual Machines have been met.

10.6.2. Thereafter Support units will be decremented from the Support Unit Allocation for work performed in response to monitoring agent alerts and/or 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 HOT FIXES IN ACCORDANCE WITH 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/OR 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
CHECK STATUS OF BACKUPS AND RESTORES	PER CUSTOMER REQUEST
PERFORM TEST RESTORES USING HOSTED BACKUP SERVICES	PER CUSTOMER REQUEST
TEST OR INITIATE REPLICATION FAILOVER FOR MANAGED REPLICA VIRTUAL MACHINES	PER CUSTOMER REQUEST
CREATION OF CUSTOMER SPECIFIC MONITORS OR REMEDIATION SCRIPTS	PER CUSTOMER REQUEST

10.6.3. Where services are (a) outside the scope of this Service Schedule (b) relate to Managed Virtual Machine exclusions, or (c) are rendered outside of Coverage Hours, or (d) are rendered 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 in accordance with Work Type and Work Roles described below:

WORK TYPES	UPLIFT
WORK OUTSIDE THE SCOPE OF THE SERVICE SCHEDULE	+40%
AFTER HOURS SUPPORT WEEKDAYS 18H00-08H00	+25%
AFTER HOURS SUPPORT WEEKENDS & PUBLIC HOLIDAYS	+50%
ESCALATION TO MICROSOFT / CITRIX / VMWARE / F5 OR OTHER VENDOR	+40%
WORK OUTSIDE THE SCOPE OF THE SERVICE SCHEDULE	+40%



WORK ROLES	UPLIFT
TIER 1 SUPPORT	-20%
APPLICATION SPECIALIST OR TIER 2 SUPPORT	+0%
INFRASTRUCTURE CONSULTANT OR TIER 3 SUPPORT	+40%
SYSTEMS ARCHITECT	+80%
PROJECT MANAGER	+40%
TIER 1 DEVELOPER	+0%
TIER 2 DEVELOPER	+40%

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

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

10.6.6. It is understood that any and all services requested of the SP that fall outside of the terms of these Managed Virtual Machines as described in this schedule will be considered projects, and will be quoted and billed as separate, individual services.

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

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

10.7.1. All Virtual Machines must be running a Microsoft Supported version of Windows Server and have all of Microsoft Service Packs and Critical Updates (released in 30 days earlier) installed.

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

10.7.3. The environment must have a currently licensed, vendor-supported hardware firewall between the internal network and the internet or an SP configure Snine firewall.

10.7.4. Fees and costs required to bring Virtual Machines up to these Minimum Standards are not included in Managed Care.

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

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

Managed Virtual Machines Exclusions

10.8. The following is expressly excluded from Managed Virtual Machines:

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

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

10.8.3. The cost to bring Virtual Machines up to minimum standards required for Services.

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

10.8.5. Programming (modification of software code) and program (software) maintenance.

10.8.6. Training Services of any kind.

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

11.1. Managed Replica VMs include replication to a second data centre in a different Fault Domain.

11.2. Replication is scheduled to run every 15 minutes, assuming the previous scheduled replication has been completed.

11.2.1. Replication is dependent on the volume of changes and cannot be guaranteed to complete within a pre-determined interval.

11.2.2. Bandwidth for replication shall be provided by SP up to a maximum of 3% (three percent) of the disk capacity included in the VM Profile in any given 12 (twelve) hour period ('**VM Profile Replication Quota**'). SP may charge extra for any bandwidth in excess of the VM Profile Replication Quota.

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

11.4. Managed Replicas consist of a pair of VM Profiles, one in each data centre.

11.4.1. **By way of example:** In order to facilitate an A3. VM Profile with replication to a second data centre, a single A3. Managed Replica will be required instead of 2 x A3. VM Profiles.



- 11.4.2. Managed Replicas requiring additional disk capacity in excess of the amount included in the VM Profile, must provision Replica Node Storage (and not locally redundant storage).
- 11.4.3. Bandwidth for replication of replica Node Storage shall be provided by SP up to a maximum of 3% (three percent) of the Replica Node Storage in any given 12 (twelve) hour period ('Replica Storage Replication Quota'). SP may charge extra for any bandwidth in excess of the Replica Storage Replication Quota.

12. Service Availability

- 12.1. If the Service is unavailable it must be reported to the SP and acknowledged by SP.
- 12.2. The period of Downtime will be calculated from when the fault is reported, SP has issued a fault report reference and has acknowledged this as a fault on the Service.
- 12.3. Following investigation and repair SP will advise the time that the Service was restored. This will be deemed to be the end of the Downtime unless the fix is not confirmed.
- 12.4. **Managed Replica Virtual Machine**
- 12.4.1. "Downtime" means the total accumulated minutes that are part of Maximum Available Minutes that have no External Connectivity.
- 12.4.2. "Monthly Uptime Percentage" is calculated using the following formula:

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

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

- 12.4.3. Exclusions to Service Availability Guarantee:
- 12.4.3.1. Any incident lasting less than 15 (fifteen) minutes.
- 12.4.4. **Service Credit:**

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