1 Complementary Services for Microsoft Azure - PaaS & Serverless

1.1 Overview

This service description covers the management of Microsoft PaaS & Serverless services and is an add-on to Managed Azure Core Services. *NTT Public Cloud Management - Managed Azure - Core Services* must also be contracted.

Managed PaaS & Serverless Services cover the following:

(a) Web Services

(b) Functions

(c) Integration Services

(d) Databases Services

(e) Analytics

The services listed below are fully supported as part of this Service Description. As Public Cloud technology evolves at such a rapid pace, it is not possible to maintain a fully up-to-date list of supported features and services. This description, therefore, serves as a baseline across common categories and may be updated from time to time by NTT.

The exact scope of the solution being delivered to the Client and related charges are clearly stated in the Statement of Work (SOW).

Category	Managed Element	Supported Services
Compute	Cloud PaaS App	. Azure App Services . Notification Hub . Azure AI Search (Cognitive Search) . Azure Logic Apps . Azure Logic Apps Custom Connector . Azure Container Apps
	Cloud Serverless Function	. Azure Function
	Cloud Virtualization	. Azure Virtual Desktop
Integration, Data & Analytics Services	Cloud Data (Integration)	. API Management . Logic Apps
	Cloud Data (Processing)	. Analysis Services . Azure Data Lake Storage . Azure Data Lake Analytics . Azure DataBrick . Azure Synapse Analytics . Azure Data Factory . Event Hub
Database Services	Cloud Database	. Azure Cosmos DB . Azure Database for MySQL . Azure SQL Database . Azure SQL Managed Instance . Table Storage . Azure Cache for Redis . Azure Database for Maira DB . Azure Database for PostgresSQL

1.2 Compute Services

(a) Cloud PaaS Apps

(i) Overview

This element of the service covers the configuration, monitoring and management of Cloud PaaS Applications. Charges are based on the number of instances present in the environment.

(ii) Supported Technologies

- Azure App Services
- Notification Hub
- Azure AI Search (Cognitive Search)
- Azure Logic Apps
- Azure Logic Apps Custom Connector
- Azure Container Apps
-

Cloud PaaS App /Web Apps	
Overview	Cloud PaaS Apps/Web Apps are HTTP-based services for hosting web applications, REST APIs, and mobile back ends
Setup Activities	. Create and set the deployment credential . Publish application; (this task can be delegated to the Client if required) . Map custom domain . Enable diagnostic logs . Upload SSL certificates . Setup deployment slots . Setup monitoring . Upload Java application . Connect to on-premise resources . Secure app . Scale app . Create a schedule backup plan
Service Request	. Stop, start, restart and delete web app . Scale app up or down . Scale instance count . Move an app to another App Service Plan . Scale an App Service Plan . Delete an App Service Plan . Modify backup schedule . Restore app service from backup . Clone an existing app . Publish application (this task can be delegated to the Client if required) . Change deployment slots (this task can be delegated to the Client if required)
Available Monitors	. Bytes Received. Bytes Sent . Cpu Percentage . Disk Queue Length . Http Queue Length . Memory Percentage. Cpu Time . Requests . Http2xx . Http3xx . Http401 . Http403 . Http404 . Http406 . Http4xx . Http5xx . Memory Working Set . Average Response Time . Average Memory WorkingSet . Http101
Service Limitation	

Azure Notification Hub	
Overview	Azure Notification Hubs provide an push engine that enables you to send notifications to most platforms (iOS, Android, Windows, etc.) from any back-end (cloud or on-premises).
Setup Activities	. Create a Namespace and Notification Hub . Configure the platform notification system (PNS)
Client Request	. Change Tier for notification hub

Sensitivity Label: General

Available Monitors	allPNSOutgoingBadOrExpiredChannelErrors allPNSOutgoingChannelErrors allPNSOutgoingExternalNotificationSystemErrors allPNSOutgoingInvalidPayloadErrors allPNSOutgoingSuccesses APNSOutgoingBadChannelErrors APNSOutgoingExpiredChannelErrors APNSOutgoingInvalidCredentialErrors APNSOutgoingInvalidNotificationSizeErrors APNSOutgoingPNSErrors APNSOutgoingSuccesses GCMOutgoingAuthenticationErrors GCMOutgoingBadChannelErrors GCMOutgoingExpiredChannelErrors GCMOutgoingInvalidCredentialsErrors GCMOutgoingInvalidNotificationFormatErrors GCMOutgoingInvalidNotificationSizeErrors GCMOutgoingPNSErrors GCMOutgoingSuccesses GCMOutgoingThrottled GCMOutgoingWrongChannelErrors	incomingScheduled incomingScheduledCanceled incomingTotal installationDelete installationGet installationPatch installationUpsert MPNSOutgoingAuthenticationErrors MPNSOutgoingBadChannelErrors MPNSOutgoingChannelDisconnectedErrors MPNSOutgoingDropped MPNSOutgoingInvalidCredentialErrors MPNSOutgoingInvalidNotificationFormatErrors MPNSOutgoingPNSErrors MPNSOutgoingSuccesses MPNSOutgoingThrottled pushNotifications registrationAll registrationCreate registrationDelete registrationGet	registrationUpdate scheduledPendingNotifications totalRequestsIncoming totalRequestsIncomingFailed WNSOutgoingAuthenticationErrors WNSOutgoingBadChannelErrors WNSOutgoingChannelDisconnected WNSOutgoingChannelThrottled WNSOutgoingDropped WNSOutgoingExpiredChannelErrors WNSOutgoingInvalidCredentialErrors WNSOutgoingInvalidNotificationFormatErrors WNSOutgoingInvalidNotificationSizeErrors WNSOutgoingInvalidTokenErrors WNSOutgoingPNSError WNSOutgoingSuccesses WNSOutgoingThrottled WNSOutgoingTokenProviderUnreachableErrors WNSOutgoingWrongTokenErrors
Service Limitation	Client is responsible for providing the certificate for authentication mode and is Out of Scope.		

Azure Logic Apps	
Overview	Azure Logic Apps is a cloud PaaS Service allows to create and run automated workflows with little to no code. Azure Logic Apps connect legacy, modern, and cutting-edge systems across cloud, on premises, and hybrid environments
Setup Activities	. Create App Service Plan and Workflow service plan . Configure Zone Redundancy. . Configure Storage (Azure Storage or SQL – Azure Storage). . Create Managed Identity. . Create Virtual Network Integration. . Create Private endpoint. . Enable App insights (Optional)
Service Request	. Event subscription to NTT approved resources. . Upload Artifacts schema, Maps and assemblies . Add Key values in Configuration . Configure Authentication with MS Identity Provider. . Modify scale up or Scale out App Service Plan. . Create connections using *Service Connector.* . Manage RBAC permissions.

Sensitivity Label: General

Service Limitation	. Create Workflows (Stateful or Stateless)
Monitors	. Workflow Runtime Logs . Function Application Logs . App Service Authentication logs (preview)

Azure Logic App Custom Connector	
Overview	Logic Apps Custom Connector allows you to register a custom HTTP endpoint with operations that will be exposed within Azure Logic Apps.
Setup Activities	. Create Logic App Custom Connector. . Associate with Integration service environment.
Service Request	. Create connector with API endpoint. . Manage RBAC permissions. . Create App Registration (OAuth Authentication)
Service Limitation	. Client has to provide Swagger file to import the configuration. . Client has to provide API definition to configure Actions, Triggers and References
Monitors	. HTTP Requests

Azure Container Apps	
Overview	Azure Container Apps is a serverless hosting service for containerized applications and microservices. Azure Container Apps enables executing application code packaged in any container and is unopinionated about runtime or programming model. Applications may scale in response to HTTP requests, events, or simply run as always-on background jobs.
Setup Activities	· Deployment source: From container image or source code or artifact · Configure container section using existing or new container image if deployment is through container image. · Configure build source using Github or artifact if deployment is through source code or artifact ·Configure Application ingress settings
Service Request	· Configure revisions and replicas · Configure scale rule setting · Configure authentication for client access to application · Create secrets to protect sensitive data like passwords and connection strings (Optional) · Configure ingress for applications that need an HTTP or TCP endpoint. · Configure custom domain for application access · Create System Assigned Identity.
Service Limitation	· A single container app environment requires at least a /23 subnet · cannot run privileged containers · Linux-based container images are required
Monitoring Alerts	

Signal	Status
Custom log search	Log Analytics
CPU Usage	Platform metrics
jvm.buffer.count	Platform metrics
jvm.buffer.memory.limit	Platform metrics
jvm.buffer.memory.usage	Platform metrics
jvm.gc.count	Platform metrics
jvm.gc.duration	Platform metrics
jvm.memory.committed	Platform metrics

jvm.memory.limit	Platform metrics
jvm.memory.total.committed	Platform metrics
jvm.memory.total.limit	Platform metrics
jvm.memory.total.used	Platform metrics
jvm.memory.used	Platform metrics
jvm.thread.count	Platform metrics
Memory Working Set Bytes	Platform metrics
Network In Bytes	Platform metrics
Network Out Bytes	Platform metrics
Replica Count	Platform metrics
Replica Restart Count	Platform metrics
Requests	Platform metrics
Reserved Cores	Platform metrics
Resiliency Connection Timeouts	Platform metrics
Resiliency Ejected Hosts	Platform metrics
Resiliency Ejections Aborted	Platform metrics
Resiliency Request Retries	Platform metrics
Resiliency Request Timeouts	Platform metrics
Resiliency Requests Pending Connection Pool	Platform metrics
Total Reserved Cores	Platform metrics
All Administrative operations	Administrative
Auth Token for Container App Dev APIs (Container App)	Administrative
Auth Token for Container App Dev APIs (Container App)	Administrative
Create or Update Container App (Container App)	Administrative
Delete Container App (Container App)	Administrative
List Container App Custom Host Name Analysis (Container App)	Administrative
List Container App Secrets (Container App)	Administrative
Start Container App (Container App)	Administrative
Stop Container App (Container App)	Administrative

(b) Cloud Virtualization

(i) Overview
This element of the service covers the configuration, monitoring and management of Azure Virtual Desktops. Charges are based on the number of workspaces. Citrix virtualization environments are not included.

(ii) Supported Technologies
- Azure Virtual Desktop

Azure Virtual Desktop	
Overview	Azure Virtual Desktop (AVD) is a cloud service that provides virtual Windows desktops and applications

Sensitivity Label: General

Setup Activities	. Create and configure virtual network and subnet . Configure security groups for target audience. . Setup firewall rules to allow connection to Session Hosts . Setup Host pool configuration – Pooled / Personal . Setup Load balancing methods for pooled host pool . Create Auto-Scale Polices – pooled host pool . Create storage for FSLogix profile setup. . Storage container configuration. . Setup FSLogix profile. . Register Desktop app group in Workspaces . Setup Application Groups – Remote / Desktop . . Configure Azure AD Connection Setup, Ignore if exists. . Setup Hybrid identity configuration, Ignore if exists. . Create User Groups. . Setup Group Policy. . Setup Multi-Factor Authentication. . Create App Attach Packages. . Setup file share storage to store App configuration . Deploy session hosts / VMs. . Create Golden Images / Custom Images . Create Manged Identity . Create Private Endpoint . Configure User access license . Setup Monitoring. . Setup backup for personal host pools . Setup patching cycles for personal host pools and OS images . Setup start/stop poilcies (Intune/GPO + on-connect start)
Service Request	. Mange User Assignment on Application Groups . Manage application groups. . Modify Group Policies . Modify Host pool configuration . Modify Auto Scale polices . Modify Session Host SKU. . Update Storage capacity . Add / Update Firewall Rules . Configure Insights. . Image build request (If the client needs to publish non-standard applications, a runbook must be provided by the client) . OS Patching . Deploy MSIX package to hosts . Modify backup setup for personal host pools
Service Limitation	. Data or existing installation migration. . Create Automation (besides scheduled start/stop) . Thrid-Party Integration. . Custom Report creation. . Setup Application Monitoring inside the Session Hosts. . VM sizing assessment . Windows License subscription when required
Monitors	. Hostpools datapoints - numberOfAvailableSessionHosts -numberOfSessionHosts . Session Hosts datapoints -sessions -status -timeFromLastHeartbeat

(c) Cloud Serverless Function

(i) Overview

This element of the service covers the configuration, monitoring and management of Cloud Serverless Function. Charges are based on the number of functions present in the environment.

(ii) Supported Technologies

- Azure Function

Azure Function	
Overview	Azure Function services run small pieces of code without application infrastructure.

Sensitivity Label: General

Setup Activities	. Create the function project and the resources needed (Resource Group, Serverless plan, and Storage account) . Publish the project to Azure . Connect Azure Services or resources to function (add bindings) . Set up continuous deployment authorization . Add deployment Azure Functions slots . Deploy function app to Kubernetes (if Kubernetes management is contracted) . Add SSL Certificate
Service Request	. Swap operations between sources and target slot . Remove a function app from Kubernetes . Enable disable a function
Available Monitors	Monitoring of Azure Function executions, data throughput and memory. . Average Memory WorkingSet: Average memory working set, in bytes . Bytes Received: Number of bytes received, per second, by the web application . Bytes Sent: Number of bytes sent, per second, by the web application . Function Execution Count: Number of executions . Function Execution Units: Number of function execution units . Http5xx: Number of HTTP 5XX responses where the server failed to fulfill an apparently valid request . Memory Working Set: Amount of memory pages touched recently by the threads in the process, in bytes
Service Limitation	

1.3 Integration, Data & Analytics Services

(a) Cloud Data (Integration)

(i) Overview

This element of the service covers the configuration, monitoring and management of Integration services. Charges are based on the number of number of services present in the environment.

(ii) Supported Technologies

- Azure API Management
- Azure App Configuration
- Azure Relay
- Azure Service Bus
- Azure Event Grid

Azure API Management	
Overview	API Management (APIM) is a way to create consistent and modern API gateways for existing back-end services.
Setup Activities	. Create an Azure API Management Service Instance . Grant access to Client to import and public APIs and use the service . Q testing . Set Api Management Policies
Client Request	. Configuration of the service for the Client to use it . QA testing . Request change policies. . Change the Api Management Limits (quotas)
Available Monitors	. Capacity . durationRaw . failedRequestsRaw . otherRequestsRaw . Status . successfulRequestsRaw . totalRequestsRaw . unauthorizedRequestsRaw . Duration . FailedRequests . OtherRequests . SuccessfulRequests . TotalRequests

Sensitivity Label: General

Azure API Management	
Service Limitation	The following are not included and are Out of Scope: . Coding of API interactions/calls or resolution/support of API is not included. . NTT does not manage the API Management environment; activities are applicable to the setup only. Management if Client responsibility.

Azure App Configuration	
Overview	Azure App Configuration centrally manages application configuration and feature flags in the cloud and integrate configuration with CI/CD processes. App Configuration complements Azure Key Vault, which is used to store application secrets.
Setup Activities	. Create App Configuration Instance(default – Free Tier) . Configure Geo-Replication. . Create Managed Identity . Create Virtual Network Integration . Create Private endpoint . Configure ARM Authentication. (default: Local) with access keys. . Configure retention for deleted stores . Configure purge protection. . Configure Private network access.
Service Request	. Change pricing tier . Regenerate Access keys – Read-write Keys / Read-only Keys . Add or Remove Key Value on configuration explorer . Add or Remove Key vault references . Create Snapshots . Archive Snapshots. . Create Feature flags. . Enable or Disable Feature Flags. . Update labels on keys in configuration store . Update labels in Feature flags . Import App configuration from exiting NTT supported App Services, App Configuration or from Configuration File. . Configure event subscription for NTT Supported resources.
Service Limitation	. Client has to provide custom configuration file to import the configuration. . Client has to provide snapshot configuration with list of keys to be included in snapshot . Client has to provide feature filter configuration details . Client has to provide label for each value in configuration explorer. . Client has to provide feature flag configuration details.
Monitors	. All HTTP Requests

Azure Relay	
Overview	Azure Relay is a cloud service that helps connect on-premises and cloud resources. It may act as a bridge between your private network and the public cloud
Setup Activities.	Create relay service Setup Hybrid connection. Setup WCF Replay connections (NetTCP or Http). Create RootManaged SAS policy. Enable or disable TLS (Transport Layer Security) versions Set up private endpoints.
Service Request	. Create Hybrid connections . Create WVF Connections. . Create Share access polices. . Rotate Account keys in Shared access Policies . Set up alerts and notifications for your Relay service
Service Limitation	. Creating Listeners and setup authentication
Monitors	. HybridConnections Events . VNet/IP Filtering Connection Logs . ListenerConnection- ServerErrors . ListenerConnection- ClientErrors

Sensitivity Label: General

Azure Relay	
	. SenderConnection – ServerErrors. . Sender Connections – Client Errors.

Azure Service Bus	
Overview	Azure Service Bus helps decouple applications and services from each other.
Setup Activities	. Create a service bus name · Create new queues or topics · Configure VNET. · Create Private Endpoints. · Create Shared Access Policies · Configure Auto Scale based on scale mode - Metric or Messaging Unit Based. · Configure Managed Identity. · Configure Disaster Recovery.
Service Request	· Create new queues or topics · Add or remove subscriptions to topics · Update access policies · Modify the Queue settings - Message Time-to-live settings, Max Size, Message lock duration. · Configure message auto-delete settings. · Create topics · Add or remove subscriptions to topics. · Enable or disable encryption · Configure auto-forwarding of messages · Configure ordering support to Topics. · Add or Remove message filters to subscriptions.
Available Monitors	· Diagnostic Error Logs · Operational Logs · VNet/IP Filtering Connection Logs · Runtime Audit Logs · Application Metrics Logs (Unused)

Azure Event Grid	
Overview	Azure Event Grid is a scalable, managed Pub Sub message distribution service that offers flexible message consumption patterns using the MQTT and HTTP protocols
Setup Activities	· Create Namespace. · Create System Topics. (Azure Service Events) · Create Custom Topics. · Create Custom Domain · Configure MQTT – Authentication (CA Client Cert / Custom JWT) · Configure MQTT - Sessions. · Create Local Authentication (Custom Events) · Configure TLS (Custom Events) · Configure event schema (Custom Events) · Configure Data Residency (Custom Events) · Create Partner registrations. · Create Partner Topics, Namespaces · Supported Partners - Auth0, MS Graph, SAP, MS EntraID · Create Managed Identity · Create Private Endpoint
Service Request	· Add or remove event subscriptions · Scale throughput units. · Change the endpoint where events are sent · Update event filtering rules · Regenerate access keys. · Add or Remove Clients from MQTT broker · Add or Remove Client Groups – MQTT broker. · Add or Remove Topic Spaces. · Update retention policy for messages in topic. · Update message Max Delivery count for subscription. · Add or Remove Delivery headers for subscription. · Grant Publisher / Subscriber permission to specific Client Group on the Topic Space. · Configure EndHub endpoint for the subscription with delivery mode set to push.

Sensitivity Label: General

Azure Event Grid	
	· Modify retry policies for failed event deliveries · Change the Event Grid domain or topic settings · Set up or modify dead-lettering for undelivered events · Add or remove access control (IAM) permissions – RBAC
Service Limitation	· The client has to provide CA Certificate to create a Client connection for the MQTT Broker. · Complex event filtering: Advanced filtering may require additional coding. · Event transformation: Changing event data before delivery may require additional coding. · Custom event schemas: handling custom event formats may require additional coding.
Monitors	. Successful Acknowledged Events . Failed Acknowledged Events . Failed Published Events . Failed Received Events . Failed Release Events . Successful MQTT Connections . Failed MQTT Connections . MQTT Disconnections . Failed MQTT Published Messages . Failed MQTT Subscription Operations . AllMetrics

(b) Cloud Data (Processing)

(i) Overview

This element of the service covers the configuration, monitoring and management of Cloud Databases. Charges are based on the number of services present in the environment.

(ii) Supported Technologies

- Azure Event Hub
- Azure Analysis Services
- Azure Stream Analytics
- Azure Data Lake Storage
- Azure Data Lake Analytic
- Azure DataBrick
- Azure Synapse Analytics
- Azure Data Factory
- Microsoft Fabric
- Azure Data Explorer Cluster

Azure Event Hub	
Overview	Azure Event Hubs is a Big Data streaming platform and event ingestion service that can receive and process millions of events per second. Event Hubs can process and store events, data, or telemetry produced by distributed software and devices.
Setup Activities	. Create NameSpace . Create Event Hub Service . Enable Event Hub Capture to store events on Blob Storage or Data Lake Storage
Client Request	. Request Management Credential for application to communicates with Event Hub. . Change Capture settings
Available Monitors	. Archive Backlog Messages . Archive Backlog Messages_Total. Archive Messages. Archive Messages_Total. Archive MessageThroughput. Archive MessageThroughput_Total. Failed Requests. Incoming Bytes Per Second. Incoming Messages. Incoming Requests. Internal Server Errors. Other Errors. Out going Bytes Per Second. Outgoing Messages. Server Busy Errors. Success full Requests

Azure Analysis Services	
Overview	Enterprise-grade analytics engine as a service
Setup Activities	. Create the Analysis Services Server and Firewall . Configure server admin users and roles . Configure Storage setting for backup . Configure the Scale-out options . Install and configure on-premises data gateway (applicable only if OS management is contracted for on-prem services)

Sensitivity Label: General

Azure Analysis Services		
Client Request	. Restore a database	
Available Monitors	ProcessingPoolIdleNonIOThreads ProcessingPoolIdleIOJobThreads ProcessingPoolBusyNonIOThreads ProcessingPoolBusyIOJobThreads MemoryThrashing MemoryLimitVertiPaq MemoryLimitLow MemoryLimitHigh MemoryLimitHard Memory MashupEngineQPU MashupEngineMemory LongParsingJobQueueLength LongParsingBusyThreads CurrentUserSessions CurrentConnections CommandPoolJobQueueLength CommandPoolIdleThreads CommandPoolBusyThreads CleanerMemoryShrinkable CleanerMemoryNonShrinkable	VertiPaqPaged VertiPaqNonPaged TotalConnectionRequests TotalConnectionFailures SuccessfullConnectionsPerSec Status ShortParsingJobQueueLength ShortParsingIdleThreads ShortParsingBusyThreads RowsWrittenPerSec RowsReadPerSec RowsConvertedPerSec QuotaBlocked Quota QueryPoolJobQueueLength QueryPoolIdleThreads QueryPoolBusyThreads QPU ProcessingPoolJobQueueLength ProcessingPoolIOJobQueueLength
Service Limitation	Managed Operating System must be contracted for on the on-premises data gateway.	

Azure Stream Analytics	
Overview	Azure Stream Analytics is a managed stream processing engine by Microsoft.
Setup Activities	. Prepare Input Data . Create Blob Storage . Define Stream Analytics Job . Configure Output Settings . Connect IoT Hub to Blob Storage . Test and Monitor . Enable Resource Logs
Client Request	. Stream Processing Configuration - Configure Input Sources - Configure Output Sinks - Monitor Job Performance . Resource Management - Modify the Number of Streaming Units (SUs) - Modify Scale Settings . Security and Access - Configure IAM roles and permissions . Data Integration - Configure New Linked Service . Networking - Configure Network Access for Integrations
Service Limitation	. Creation of Stream Processing Jobs . Define Stream Processing Queries . Write SQL Queries . Define Stream Processing Queries . Limitation on the Number of Streaming Units (SUs) . Restrictions on Configuring IAM Roles and Permissions . Limits on Integrating with Linked Services and Data Sources . Constraints on Network Access Configuration . Limitations on Monitoring Job Performance . Creation of pipelines, Power Query, Data Flow Functions and Data Sets, Ingestion of data and Data transformation activities . Create Spark Notebooks or Spark Definitions . Create or Modify Triggers

Sensitivity Label: General

Data Lake Storage	
Overview	Scalable, secure data lake for high-performance analytics
Setup Activities	. Create Storage Account . Transfer Data when source data is under management or set permission for Client can transfer data . Provide Access to the Client to manage the data.
Client Request	. Change Access Control
Available Monitors	. DataRead_Raw . DataWritten_Raw . Read Requests_Raw . StatusCode . TotalStorage_Raw . WriteRequests_Raw
Service Limitation	

Data Lake Analytics	
Overview	Data Lake Analytics is an on-demand analytics job service that simplifies big data.
Setup Activities	. Create Data Lake Analytics Account . Add data sources . Add new users or security groups to the Data Lake Analytics account . Grant access to the Client to manage jobs, and monitor pipelines jobs . Apply account polices
Client Request	. Add/ Remove data sources . Adjust quotas and limits for Azure Data Lake Analytics
Available Monitors	. JobAUEndedCancelledRaw . JobAUEndedFailureRaw . JobAUEndedSuccessRaw . JobEndedCancelledRaw . JobEndedFailureRaw . JobEndedSuccessRaw . Status
Service Limitation	Troubleshooting of U-SQL jobs is not included and is Out of Scope.

Azure DataBricks	
Overview	Azure DataBricks is a data analytics platform optimized for the Microsoft Azure cloud services platform. Fast, easy, and collaborative Apache Spark-based analytics platform).
Setup Activities	. Create the Azure DataBricks Workspace and the resources needed (Resource Group and Pricing Tier) . Deploy Azure DataBricks workspace into existing Virtual Network.. Create Azure DataBricks Cluster – Choose Cluster mode, DataBricks Runtime Version, Worker Types, Driver Types etc. Configure RBAC Guardrails.. Ensure LogicMonitor monitoring for Worker Nodes is in place. Ensure Cloud Health Billing is in place for cost management.
Client Request	. Terminate, Restart or Delete Cluster. Set Permission settings on Cluster.. Change Node Instance Type. Change DataBricks Runtime Version. Configure Cluster Access Restrictions. RBAC Permission changes. Opening of Microsoft Support requests
Available Monitors	. CPU Utilization: . Network Throughput. . Disk IOPS . Disk Throughput . Resource Health . Virtual Machine Resource Health
Service Limitation	The following are not included and are Out of Scope: . Creation of new Notebooks . Creation of Tables.

Sensitivity Label: General

Azure DataBricks	
	. Creation of new Jobs . Creation of new MLFlow Experiments. . Library imports.

Azure Synapse Analytics	
Overview	Azure Synapse is an enterprise analytics service that assist with time to insight across data warehouses and big data systems.
Setup Activities	. Create Synapse Workspace . Create Data Lake Storage Gen2 . Configure EntraID and SQL Server local Authentication. . Configure Managed VNET . Create Private Endpoints . Configure Storage Blob Data Contributor Access for Workspace Identity on Data Lake. . Configure Managed Identities (if required) . Configure Private or Public Self-hosted integration runtime connectivity . Configure New linked service (ie Azure SQL Database/Storage Account) for services that NTT manage
Client Request	. Create SQL Pools . Create Apache Spark Pools . Modify the Number of Nodes in Spark pools . Modify Scale Settings for Apache Spark Pool . Modify the Spark pool Intelligent Cache Size . Configure Automatic Pausing for Apache Spark Pool . Modify SQL Dedicated Pool Performance Level . Reset SQL local Credentials . Configure IAM roles and permissions for client access. . Configure New linked service (ie Azure SQL Database/Storage Account) for services that NTT manages . Configuration of any network access required for integrations . Git Configuration with Client Azure DevOps / GitHub Repos
Service Limitation	. Creation of pipelines, Power Query, Data Flow Functions and Data Sets, Ingestion of data and Data transformation activities . Create Spark Notebooks or Spark Definitions . Create or Modify Triggers

Azure Data Factory	
Overview	Hybrid data integration at enterprise scale, made easy
Setup Activities	. Create Azure Data Factory v2 · Configure Git repository with either Azure DevOps or GitHub using Client provided details. · Enable or Disable Managed virtual network setting · Configure Private Endpoints · Configure Private or Public Self-hosted integration runtime connectivity · Enable Data factory Encryption for Client owned keys (requires Azure Keyvault) - Optional · Assign Tagging Policy . Configure IAM roles and permissions for Client access. · Configure Managed Identities (if required) · Configuration of any network access required for integrations. · Configure New linked service (ie Azure SQL Database/Storage Account) for services that NTT manage via Data Factory Studio
Client Request	. Enable Data factory Encryption for Client owned keys (requires Azure Keyvault) - Optional · Configure IAM roles and permissions for client access. · Configure Managed Identities (if required) · Configuration of any network access required for integrations. · Configure New linked service (ie Azure SQL Database/Storage Account) for services that NTT manage via Data Factory Studio
Available Monitors	. activityCancelledRuns . activityFailedRuns

Sensitivity Label: General

Azure Data Factory	
	. activitySucceededRuns . factorySizeInGbUnits . factorySizeInGbUnitsMaxAllowed . integrationRuntimeAvailableMemory . integrationRuntimeAvailableNodeNumber . integrationRuntimeAverageTaskPickupDelay . integrationRuntimeCPUPercentage . integrationRuntimeQueueLength . pipelineCancelledRuns . pipelineFailedRuns . pipelineSucceededRuns . resourceCount . resourceCountMaxAllowed . triggerCancelledRuns . triggerFailedRuns . triggerSucceededRuns
Service Limitation	The following are not included and are Out of Scope: . Creation of pipelines, Power Query, Data Flows and Data Sets, Ingestion of data, Transformation of data, Configuration of SSIS packages · Data transformation activities

Microsoft Fabric	
Service Overview	Microsoft Fabric is an end-to-end analytics and data platform for enterprises that require a unified solution. It encompasses data movement, processing, ingestion, transformation, near real-time event routing, and report building.
Setup Activities	. Create Fabric Tenant . Create Fabric Capacity . Configure Private Link Service . Configure VNET . Configure VNET Peering to Hub. . Create Managed Private Endpoints for fabric Notebooks. . Enable Disaster Recovery for OneLake (supported only with F64 or above plans) . Configure Azure Data Factory Integration with Fabric (Optional) . Configure Private or Public Self-hosted integration runtime connectivity to Azure Data factory (Optional) . Create Workspace from Fabric Capacity. . Configure Git Integration. . Create Workspace Identity. (supported only with F64 or above plans) . Configure Trusted Workspace access to ADLS Gen2. . Configure OneLake Shortcut to Storage Account with Trusted workspace access
Service Requests	. Create Fabric Capacity . Resize Fabric Capacity . Create new Workspace . Manage users permission in OneLake RBAC and Workspace Roles. . Configure Fabric Capacity Notifications. . Create Workspace Identity . Create OneLake Shortcuts for data services that NTT manages with Trusted Workspace Access. . Configure Trusted workspace access to ADLS Gen2 . Create Managed Private Endpoints for Notebooks. . Create Private Endpoints and connect it to Fabric Private Link Service to access Fabric Items . Git Configuration with Client Azure DevOps / GitHub Repos.
Service Limitations	. Creation of task flows, pipelines, Power Query, Data Flow Functions and Lake houses, Ingestion of data and Data transformation activities . Create Spark Notebooks or Spark Definitions. . Create or Modify Triggers. . Configure data hubs, Direct Lakes and Score models

Azure Data Explorer Cluster	
Overview	Azure Data Explorer is a managed analytics platform that makes it assist with analyzing high volumes of data in near real time.

Sensitivity Label: General

Azure Data Explorer Cluster	
Setup Activities	· Create Compute instance – Dev/Test, Storage Optimized and Compute Optimized · Configure Availability Zone · Configure default instance count or configure optimized auto scale (Storage / Compute Optimized instance only) · Configure cluster auto-stop. · Create Managed Identity · Create Virtual Network Integration · Create Private endpoint · Configure Double encryption and Disk encryption. · Configure Retention and cache period for Database.
Service Request	· Scale up or Scale out Compute size. · Enable language extension grid for R or Python Language. · Upload Python custom image. · Manage RBAC on cluster. · Add or Remove User assigned Identity. · Configure security settings for cluster. · modify tenant permissions. · Create Data ingestion connection for NTT supported Resources. · Manage RBAC permissions on Databases. · Manage permissions on Data Ingestion URI
Service Limitation	· Create table in Database · Define Database schema · Configure Streaming Ingestion, Enable Purge · Create Database · Create custom KQL Queries for Visualization. · All ETL activities in Databases or query databases. · Create Rest API for Data management.
Monitors	· Succeeded Ingestion · Failed Ingestion · Ingestion Batching · Command · Query · Table usage statistics · Table details

Azure Data Explorer Cluster	
Overview	Azure Data Explorer is a managed analytics platform that makes it assist with analyzing high volumes of data in near real time.
Setup Activities	· Create Compute instance – Dev/Test, Storage Optimized and Compute Optimized · Configure Availability Zone · Configure default instance count or configure optimized auto scale (Storage / Compute Optimized instance only) · Configure Streaming Ingestion, Enable Purge. · Configure cluster auto-stop. · Create Manged Identity · Create Virtual Network Integration · Create Private endpoint · Configure Double encryption and Disk encryption. · Create Database · Configure Retention and cache period for Database.
Service Request	· Scale up or Scale out Compute size. · Enable language extension grid for R or Python Language. · Upload Python custom image. · Manage RBAC on cluster. · Add or Remove User assigned Identity. · Configure security settings for cluster. · modify tenant permissions. · Create Data ingestion connection for NTT supported Resources. · Manage RBAC permissions on Databases. · Manage permissions on Data Ingestion URI
Service Limitation	· Create table in Database · Define Database schema · Create custom KQL Queries for Visualization. · All ETL activities in Databases or query databases.

Sensitivity Label: General

Azure Data Explorer Cluster	
	· Create Rest API for Data management.
Monitors	· Succeeded Ingestion · Failed Ingestion · Ingestion Batching · Command · Query · Table usage statistics · Table details

1.4 Database Services

(a) Cloud Database

(i) Overview

This element of the service covers the configuration, monitoring and management of Cloud Databases. Charges are based on the number of instances present in the environment.

(ii) Supported Technologies

- Azure Cosmos DB
- Azure Cache for Redis
- Azure SQL
- Azure DB For MySQL
- Azure DB for PostgreSQL

Azure Cosmos DB	
Overview	Azure Cosmos DB is a fast NoSQL database with open APIs for any scale.
Setup Activities	. Creation of database account and settings . Creation and addition of database and collection . QA testing . Configuration of consistency level . Configuration of CosmosDB firewall security policies . Delivery of connection string keys to the Client . Enabling turnkey global distribution . Configuration of CosmosDB replication (if applicable) and failover priorities
Service Request	. Add/remove regions from database distribution . Create/update/delete Cosmos DB accounts . Change tags . Regenerate account key . Modify failover priority . Create/delete resources, databases, collections and documents
Available Monitors	. Metadata Requests: Number of metadata requests since the last polling interval . Mongo Request Charge: Number of Mongo request units being consumed . Mongo Requests: Number of Mongo request made since the last polling interval . Total Requests: Total number of requests made since the last polling interval . Total Request Units: Total number of request units being consumed . Available Storage: This metric will take your available storage about the aggregate size from each collection data and index
Service Limitation	

Azure Cache for Redis	
Overview	Azure Cache for Redis provides an in-memory data store based on open-source Redis software.
Setup Activities	. Configuration of the FQDN server name . Configuration of instance type . Delivery of connection string key to the Client . QA testing . Enabling cache diagnostics . Configure memory policies . Configure updates schedule If Premium instance is contracted:

Azure Cache for Redis	
	. Configuration of clusters . Configuration of data persistence . Configuration of virtual network . Configuration of scheduled updates
Service Request	. Change basic parameters like name, ports, pricing tier and selected cache metrics . Add/change/remove tags . Regenerate access keys . Enable/disable keyspace notifications . change updates schedule . Import/export data . Reboot services
Available Monitors	. Connected Clients: Current number of connected clients . TotalCommandsProcessed: Total number of commands processed . Cache Hits: Total number of cache hits . Cache Misses: Total number of cache misses . Cache Read: Number of bytes/sec read from cache . Cache Write: Number of bytes/sec written to cache . Evicted Keys: The number of items evicted from the cache during the specified reporting interval due to the maximum memory limit . Expired Keys: The number of items expired from the cache . Get Commands: The number of GET operations from the cache . CPU: Percent CPU utilization.High CPU . Server Load: The percentage of cycles in which the Redis server is busy processing and not waiting idle for messages . Set Commands: The number of SET operations from the cache . TotalKeys: Total number of keys . Used Memory: Amount of memory used, in bytes . Used Memory Rss: Amount of RSS (Resident Set Size) memory used, in bytes for Cluster Configurations extra monitors are available: . Cache Hits: Total number of cache hits . Cache Misses: Total number of cache misses. . Cache Read: Number of bytes/sec read from cache. . Cache Write: Number of bytes/sec written to cache. . Connected Clients: Current number of connected clients. . Evicted Keys: The number of items evicted from the cache during the specified reporting interval due to the maximum memory limit. . Expired Keys: The number of items expired from the cache. . Get Commands: The number of GET operations from the cache. . CPU: Percent CPU utilized by the specified shard. . Server Load: The percentage of cycles in which the Redis server is busy processing and not waiting idle for messages. . Set Commands: The number of SET operations from the cache. . Total Commands Processed: Total number of commands processed . . Total Keys: Total number of keys. . Used Memory: Amount of memory used, in bytes. . Used Memory Rss:Amount of RSS (Resident Set Size) memory used, in bytes.RSS memory is the amount of space used by the memory pages consumed by the process. . Cache Hit Ratio: Percent cache hit ratio / total Cache . Cache Miss Ratio: Percent cache miss ratio/ Total Cache
Service Limitation	. Scale down (from higher tier to lower tier) is not possible and is Out of Scope

Azure SQL	
Overview	Azure SQL Database is a fully managed Platform as a Service (PaaS) Database engine
Setup Activities	. Create Azure SQL instance . Creation and configuration of the users/databases/permissions required . Create server-level firewall rules . Database configuration (static or elastic) . Data import using native Azure SQL tools . Configuration of database geo-replication . Configuration of elastic pool

Sensitivity Label: General

Azure SQL	
	. Create a master key (if applicable) . Create or obtain a certificate protected by the master key (if applicable) . Create a database encryption key and protect it by certificate (if applicable) . Set the database to use encryption (if applicable)
Service Request	. Change tier size . Data import and export using native Azure SQL tools . Management of elastic pool settings . Management of database users and permissions . Management of security and server-level firewall rules . Changes to SQL geo-distribution, activation of geo-replicated failover database (if applicable) . Restore to any point in time within the retention period . Execution of End User provided scripts . Turn on/off SQL Database Auditing . Enable/disable SQL database threat detection notification . Enable Always Encrypted feature
Available Monitors	. Batch Req: . Buffer Cache Hit Ratio: . Buffer Cache Hit Ratio Base: . Compilations: . Connections: . Deadlocks . Full Scans . Index Searches . Latch Waits . Latch Wait Time . Lock Timeouts . Lock Waits . Lock Wait Time . Logouts Sec . Page Life . Page Reads . Page Splits . Page Writes . Recompilations . Status Transactions . Work Files . Work Tables . Buffer Cache Hit Ratio Perc . New Connection Ratio . Page Life Non Zero
Service Limitation	

Azure DB for MySQL	
Overview	Azure Database for MySQL is a relational database service based on MySQL Community Edition
Setup Activities	. Create a MySQL single server . Configure a server-level firewall rules . Configure SSL connectivity . Configure server parameters . Create databases and users . Create the service endpoint or Private Link . Configure local or geo-redundant backup . Configure Data-in replication (for hybrids solutions)
Service Request	. Restore a database (point- in-time or Geo-restore) . Create a read replica
Available Monitors	. Active Connections . Compute Consumption Percent . Compute Limit . Connections Failed**Legal / Service Descriptions** . CPU Percent . IO Consumption Percent . Memory Percent

Sensitivity Label: General

Azure DB for MySQL	
	. Storage Limit . Storage Percent . Storage Used
Service Limitation	

Azure DB for PostgreSQL	
Overview	Azure Database for PostgreSQL is a managed relational database service. It is based on the open-source PostgreSQL database engine.
Setup Activities	. Determine the Azure region closer to users . Create Flexible Server (production-use) or Single Server (test or development use) . Create Server with latest PostgreSQL version (or previous version according to requirements) . Configure basic settings such as unique server name, workload type . Configure compute and storage based on performance requirements . Configure Private Network (uses delegated subnet) and/or Public Network access (based on requirements) . Setup firewall rules to allow access from specific IP address or ranges . Setup admin access (default is PostgreSQL authentication) . Utilize tools like psql or preferred PostgreSQL client to connect . Create databases according to requirements (default is postgres) . Configure autogrow . Implement indexing and query . Configure High availability (as per requirements) . Configure backup and retention policies. Configuration Encryption
Service Request	. Modify compute and storage options based on performance and cost . Create additional databases . Enhance performance by doing proper indexing and query optimization . Modify Backup and retention policies . Maintain and Modify firewall rules for known IP address(s) or range(s) . Modify encryption at reset and in transit . Modify maintenance schedule (if different than system-managed) . Plan and perform major version upgrade . Manage Stop, Start and Restart of Server . Manual backup supported using pg_dump
Service Limitation	. Creation of tablespace is unsupported . Each pricing tier has default maximum number of connections by vCore configuration . Server storage can only be scaled in 2x increments . Maximum storage size for Flexible Server type is 32 TB . Decreasing server storage size is currently unsupported . Moving in and out of a virtual network is currently unsupported . Combining public access with deployment in a virtual network . Postgres 10 and older versions are unsupported . Manual backup from Azure . Stopped Azure Database for PostgreSQL flexible server instance automatically starts after 7 days . Change in Custom maintenance window takes effect only with the following monthly scheduled maintenance

Sensitivity Label: General