CORE

TRANSFORMING BUSINESS
THROUGH TECHNOLOGY

AURORA IDAM

SERVICE DESCRIPTION

Contents

1. Management Summary

The purpose of this document is to overview the services offered and included in Core's IDAM as a Service offering – Aurora. Aurora is a fully operable IDAM solution that matches or exceeds the capabilities of other global software vendor provided IDAM offerings.

Aurora has some unique features in the overall service offering;

- Aurora is a decentralised IDAM service, deployed directly and self contained in the customers own IT environment hosted in Microsoft Azure. Its decentralised nature ensures that all identity data and attributes only remain in the customers own operational perimeter.
- Aurora is based on native Microsoft technologies, ensuring that all underlying system components are evergreen, constantly evolving in line with overall technology development in Microsoft Azure.
- The architecture of Aurora ensures that the customer always owns user identities, any future removal of Aurora will only remove the specific automation and synchronisation functionality that Aurora provides, existing identities will not be negatively affected.
- Aurora is designed to compliment or augment existing standard Microsoft identity services, such as Active Directory and Entra ID, but provides extensibility to support identities and authentication across 2,500+ technology platforms
- Aurora provides a wide range of automation and synchronisation capabilities for a customer, with key features such as;
    - Automated JML processes through integration with relevant authoritative sources such as HR or ITSM systems, contractor management systems, or services such as Azure / Entra ID B2B.
    - User license assignment and management as part of the provisioning and deprovisioning process.
    - Password synchronisation across all connected services, ensuring simple access for users to all in scope platforms.
    - Identity attribute transformation services to enable interoperability of users with a wide range of platforms including legacy on premises line of business systems that may require specific user identity structures.
    - Allocation of an Immutable ID to all user identities, which enables continuity of user identity and audit in any scenario where user attributes such as name or employment status may change and impact the UPN of that user.
    - Native PIM and PAM management, including the ability to provision parallel privileged accounts to relevant users as part of the automated JML process, just in time account activation and time boxed access with session specific authentication.
    - One touch deprovisioning of users which will immediately suspend any access rights of a user, with automated deprecation and clean up activities

2. Description and Purpose

The Aurora IDAM platform provides a policy-based administration solution that meets the needs of modern enterprises.

The administrative policy enforcement featured by Aurora IDAM considerably reduces administrative workload, improves network security, and ensures consistency across the entire enterprise. Automating administrative workflow significantly reduces the amount of time to complete tasks and can eliminate certain tasks altogether. It also minimizes errors, reduces the need for rework, and combines related actions into a single batch. Aurora IDAM provides the facility to specify how, when, and what must change, whenever directory objects are created, modified, or deleted.

Furthermore, it is possible for Aurora IDAM to only accept data changes that conform to certain formatting requirements. This helps maintain control of the data stored in the directory as well as how it looks in other systems like email. For example, when creating a user account for a new employee, Aurora IDAM can automatically retrieve information from Active Directory via the Core Integration Engine, use it as the default information in the user account properties, create a home folder and home share, and add the new user to security groups.

This entire procedure equates to one task, but without Aurora IDAM, it could be ten or more. With the ability to enforce administrative policies and automate administrative workflow, Aurora IDAM not only saves time, but also keeps network objects in a consistent state in relations to each defined policy. This addresses important security, usability, and integrity issues that are central to the management of network object data.

In Aurora IDAM, administrative policies are defined by using Policy Objects—collections of policies. Policy Objects define the behaviour of the system when directory objects are created, modified, or deleted. There is the ability to create a Policy Object that includes any number of different policies, such as format validation, generation rules for the values of object attributes, scripts that supplement administrative operations, automatic creation of user mailboxes on prescribed Exchange servers, automatic creation of user home folders and home shares, and relocation of an object to a specified container when it meets certain criteria such as changing Department.

Aurora IDAM provides extensive capabilities for automating administrative processes. Policy Objects can run customizable scripts before or after the execution of any specific task, and multiple tasks can be combined into one operation. This functionality significantly reduces the amount of time to complete administrative tasks and minimizes errors.

As well as Policies, Aurora IDAM also provides a rich workflow system for directory data management automation and integration. Based on Microsoft's Windows Workflow Foundation technology, this workflow system enables IT to define, automate and enforce management rules quickly and easily.

Workflows extend the capabilities of Aurora IDAM by delivering a framework that enables combining versatile management rules such as provisioning and de-provisioning of identity information in the directory, enforcement of policy rules on changes to identity data, routing data changes for approval, e-mail notifications of particular events and conditions, as well as the ability to implement custom actions using script technologies such as Microsoft Windows PowerShell. Workflows may also include approval rules that require certain changes to be authorized by designated persons (approvers). When designing an approval workflow, the administrator specifies which kind of operation causes the workflow to start and adds approval rules to the workflow. The approval rules determine who is

authorized to approve the operation, the required sequence of approvals, and who needs to be notified of approval tasks or decisions. The workflow can also cause automatic escalation if an approval sits waiting for more than a number of days, designated secondary approvers can automatically be emailed to approve the change so the request doesn't go stagnant in the IT system.

By delivering e-mail notifications, workflows extend the reach of management process automation throughout the enterprise. Notification activities in a workflow let people be notified via e-mail about events, conditions or tasks awaiting their attention. For example, approval rules can notify of change requests pending approval, or separate notification rules can be applied to inform about data changes in the directory. Notification messages include all necessary supporting information, and five provide hyperlinks enabling message recipients to take actions using a standard Web browser, or to accept or reject the change via email (if using a supported email system).

The Aurora platform integrates with (but does not include by default), Windows Azure Active Directory Premium which then offers Single Sign On to these cloud-based services plus around 2,500 other cloud applications

2.1. Synchronisation

Aurora iDAM is configured to create and manage Identities using approved authoritative sources. These can include HR systems, employee databases, contractor databases or other Directory services such as a supplier Active Directory / Entra ID environment or Entra ID B2B.

The created identity is synchronised to a "Principal Directory Service" that is hosted in the customers environment. Typically this will be Entra ID or a hybrid Active Directory instance.

2.2. Transformation

Core completes a wide range of necessary attribute transformations in order to enable a cloud identity to work in legacy environments where required by the customer.
These transformations are completed "on the fly" by Aurora, which uses proprietary code and mapping tables to amend, alter, correct, or re-organise cloud attributes in the Aurora engine, depositing the transformed identity into relevant legacy domain active directory.
The transformations use multiple sources of information;

- Information captured from the authoritative source systems to provide known good data for elements such as organisational hierarchy
- Mapping tables to correct and standardise free type elements such as office locations in transformed identities.

This is not an exhaustive list, but it provides a view on the types of changes being made by the Aurora platform, where the customer has defined requirements.

2.3. Immutable ID

All identities processed by Aurora IDAM will be assigned an “Immutable ID”. This attribute is designed to provide a permanent anchor for the users’ specific Identity that will allow it to be recognisable and linked to prior versions of itself as a result of the following changes;

- Changes to the UPN (User Principle Name) that may happen as a result of a number of factors related to the individual user
- Changes to the employment status of the user, including change in temporary or permanent employment, i.e. contractor to permanent employment, or in larger organisations, migration from one of the customers group entities to another.

2.4. Other Identity related services

Aurora IDAM can also be used to provide other services relating to the corporate identity of users to services that require them via a reporting extract that can be automatically delivered to consuming services.

Typical services include;

- HR Risk Management or tracking systems
- Room and Desk Booking solutions
- Physical Security Systems

2.5. Microsoft License Assignment

Aurora IDAM manages the provisioning of Microsoft licensing (including bolt on components) to users based on their role and group memberships as part of the processing of Joiners, Movers and Leavers activity. This functionality can be extended to non Microsoft licenses where required.

Once a new identity is provisioned, or a user changes roles or moves organisation, in Aurora, licenses are assigned via workflow to the user based on their role type, which includes the enabling of specific sub features of these M365 licenses as part of the process.

2.6. Aurora Production

Aurora IDAM includes production and Pre Production versions as part of the overall offering. Production is the live version of Aurora that is used to deliver operational changes to enable business as usual activity under the provision of Identity services. This is hosted and owned by the customer, but supported and managed by Core.

The service is scalable from 1,000 to 250,000+ users within agreed performance parameters. The service can be scaled up or scaled out to deliver consistent performance if the number of users under management increases above the current size threshold, although it should be noted that this will incur additional hosting fees for larger or higher volume of IaaS.

2.7. Aurora Pre-Production (PPE)

Aurora IDAM also contains a Pre-Production environment which is designed and available to enable testing of proposed changes to identity scope or activities in a sandboxed environment that will cause no operational impacts in the event of an error, omission, or unexpected outcome from developmental projects.

The Pre-Production environment is scaled to support 1/10th users at a slightly reduced performance standard from the production environment. The environment also contains facsimiles of any connected domains or services in use by the customer.

These domains are configured to be as close to the configuration of the production domains they represent in legacy as is possible in an Azure environment (there are limits to the OS level and functional levels that can be provisioned in Azure, so some of the PPE domains are hosted on newer versions of the OS).

PPE is typically not populated with users when not in use. When required, Aurora PPE is populated with a sub section of pseudonymised users from the live production environment to enable testing.

2.8. Principal Directory Support

Core's service also provides support for some elements of the Principal Domain where Identities are provisioned by Aurora.

Typically the support relates to the patching, updating and overall management of the domain controllers, and support related to Entra ID Connect / Cloudsync where hybrids exist.

3. Technical Information

3.1. Azure Datacentres and Connectivity

The Aurora IDAM service is designed to be hosted upon Microsoft's Azure platform, ideally inside the customers existing environment, or in another customer owned Azure instance. One of the fundamental principles of our offering is that all activity takes place inside your own security perimeter.

All administration is carried out with separate, user identifiable accounts to ensure clarity of auditing and access is provided using "Least Privilege" principles.

Availability

Below are further details regarding the major components within the IDAM solution.

- For all Virtual Machines that have two or more instances deployed in the same Availability Set, the solution will have external connectivity at least 99.95% of the time.
- For Cloud Services, when two or more role instances (Virtual Machines) are deployed in the same domains, the Internet facing roles will have external connectivity at least 99.95% of the time.
- Aurora IDAM services are setup with Geo Redundant Storage (GRS). Process request data will be successfully read (and written) from (and to) GRS Accounts at least 99.9% of the time.
- Each VPN Gateway is guaranteed to have 99.9% availability.
- DNS queries will receive a valid response from at least one of our Azure Traffic Manager name server clusters at least 99.99% of the time.

The High Availability (HA) model for the Aurora solution is based on vendor best practice for HA for the software in use, and best practice from Microsoft for the Azure platform. We have designed the solution to achieve the Microsoft stated virtual machine SLA of 99.95%.

From a Recovery Time Objective (RTO) and Recovery Point Objective (RPO) we look at the different layers of the solution.

- Aurora V3 utilises multiple backup features. At the storage level, Aurora holds;
    - Daily backups for each of the 7 preceding service days
    - 1 monthly backup for the preceding service month
    - 1 6 monthly backup

- Primary focus for any DR event is the last good daily backup. Because of the configuration of Aurora, which looks for deltas between data sets, the RPO for storage is low as long as the wider DEFRA identity estate is available and intact.
    - The RPO for storage is stated as 0 minutes.
    - The RTO for storage is stated as 5 minutes.

- At the virtual machine layer, we deploy the virtual machines in pairs inside Availability Sets. The availability set uses anti-affinity to guarantee the virtual machines are not running inside the same racks. An outage of a whole rack which may affect the storage or virtual machine layer can only take out a single instance and traffic will be directed to the surviving node ensuring a near instant RTO in the case of rack failure.

  The RTO of the node that suffers the failure, or a node deployed without HA, is the reboot time of the Windows operating system (as it will automatically boot up in one of the 5 other locations) as well as the startup time of the application services that run on the server.
    - The RPO for Virtual Machines is that of the storage sublayer where the Virtual Disks are stored.
    - The RTO for Virtual Machines is stated as 10 minutes.

- Virtual Machines are also backed up under a schedule in the event that we need to restore following a major event that affects availability sets. These are completed on the following schedule;

    - Daily backups for each of the 7 preceding service days
    - Weekly backups covering the 4 preceding service weeks
    - Monthly backups covering the 6 preceding service months

- At the networking layer, we use Traffic Manager to direct client requests to the various cloud service publishing end points. The traffic manager settings have been configured for a 300 second Time To Live (TTL) setting. The TTL tells the client machine how long it should cache the IP address that Traffic Manager gives it. When a client machine connects to an Aurora service the Traffic Manager will give the IP address of a running node. If that node then fails Traffic Manager stops issuing that IP address to new requests. If, however, a client connects to an Aurora service which then fails, the client will attempt to use that IP address until the Time to Live expires which is up to 5 minutes.

    - The RPO for Traffic Manager is stated as 0 minutes as no data is stored in Traffic Manager.
    - The RTO for Traffic Manager is stated as 5 minutes as this is the configure Time To Live of the service.

- Because the availability of the service is layered, the RTO is based on the recovery time of each component in a worst case scenario.
  Storage RTO + Virtual Machines RTO + Traffic Manager RTO
  0 + 10 min + 5 min = Overall RTO of 15 minutes.

- Core do not guarantee that each client will recover access to the service within the RTO timeframe of 15 minutes as Internet based components can interfere with the Time To Live on DNS. It is a known issue with DNS that the Time To Live stated on a single DNS record is the TTL that the owner would like the client to honour. However, DNS caching servers in between the client PC and the DNS record managed by Traffic Manager may have their own cache time-outs which are out of the control of Core and the client operating system.

  This means in the event of a failover the client computer may purge the IP address of the dead node at the 5-minute mark, start a new DNS request and still be passed the old IP of the dead node by an intermediary DNS server, possibly the clients own internal DNS server,

of the upstream DNS server of the client's Internet Service Provider.

- The overall availability of the platform is stated as the time taken for the platform to be visible back to the internet in general. Individual computers, dependent on their physical location and network setup may not achieve the stated RTO and Core do not take this as a failure of the platform as we do not control the internet and services between Azure and the client.

- The RTO of the Core Integration Engine synchronisation service is dependent on access and recovery times back to the source systems from which we sync. As these will vary by connection, we do not state an RTO on the synchronisation service.

Disaster Recovery Drill Schedule

In Aurora, Core provide Disaster Recovery testing as part of the baseline service.

A partial Disaster Recovery (DR) drill will be conducted monthly in the PPE (Pre-Production Environment) during the first weekend of each month. This testing is strictly limited to PPE and is designed to validate the failover process, ensure recovery times are within acceptable limits, and confirm that all dependencies function as expected.
In addition, a partial Production DR test will be conducted once per year. This annual exercise is carefully planned to validate recovery procedures in the live environment with minimal risk and impact.

Conducting monthly PPE-based DR drills and an annual Production DR test helps to:

- Identify and address potential issues before an actual disaster occurs.
- Validate restore and recovery procedures in both PPE and Production contexts.
- Ensure that all stakeholders are familiar with the recovery process.
- Minimize downtime and data loss during real incidents.
- Strengthen operational resilience through ongoing testing and refinement.

By maintaining this schedule, we ensure that the Aurora environment remains resilient, well-prepared, and capable of handling unexpected disruptions with minimal impact.

3.2. Scope of service wrapper and responsibilities

Functions	Included within the service	Client responsibilities	Out of scope and chargeable
Service Operations • Service Desk • Incident management • Security Incident Management • Problem Management • Escalations	24x7 3rd line support and incident escalation to the client via their ITSM helpdesk via email. Documentation and knowledge sharing of known issues (This includes defined and pre-installed applications, encryption configuration, policy and security settings)	All user-based escalations to be triaged via the client. 1st line support, with assistance from Core Technology Systems during setup phase. 1st line support from client's service desk	Any enhancements to service offerings such as further support requirements
Service Operations • Event Management • Technical Operations and monitoring • Request Fulfilment • Access Management	• Patch Management including all software deployed in the base service build • Monitoring of patch/security update installation success • 24\7 monitoring of system and associated components • 3rd line support of service requests • Notification of security incidents and breaches. The notification is dependent on level and circumstance.	Management and co-ordination with 3rd parties Contributing to Change Management processes required for patching and infrastructure changes as required on the platform. 1st line triage of service requests and reactive incidents (	Additional client customisation e.g. additional client on-boarding, application on-boarding. Event forwarding to systems outside Core's management backend. E.g. Client SIEM or ticketing system.
Service Transition Functions • Change Control • Release Management and deployment • Asset and configuration management	Provision of agreed change impact review and control process in line with Core Technology Systems service standard. Provision of process for staged release of all solution deployments, software updates, policy changes and patches.	Co-ordinating Change Management contributions from 3rd parties where management of components is outside of Core's scope. User Acceptance Testing of changes to the solution and sign off	Changes to the solution that are required from 3rd party components and\or resulting in enhancements outside of the agreed solution
Continuous Service Improvement	Monthly service review meetings as agreed and monthly performance reporting Providing updates and agreed enhancements that are provided by the solution providers.	Managing 3rd party components that require changes and enhancements to enable service improvement initiatives	Service improvement initiatives required from 3rd party components

3.3. Underlying Technical Services

The following underlying technical services are used to deliver this service. For clarity, Core is acting as the Prime Contractor for the delivery of the services in the table below delivered by Microsoft , all other services Core is simply a service provider.

Service	Service Provider	Service Type
Microsoft Azure IaaS	Microsoft	Hosted
Microsoft Windows Federation Service	Microsoft	Hosted
Core Integration Engine	Core	Hosted

3.4. Service security, availability, patching and continuity

All Core managed servers are part of a 2 tier patching regime.

All odd numbered servers patch on Tuesdays to catch the 'patch Tuesday' patch set from Microsoft which is the second Tuesday of the month. All even numbered servers patch on Sunday to allow a sufficient amount of time to pass to block any unstable patches that may adversely affect our managed environment. This regime ensures that only half of the estate will be affected in the case of a bad patch being deployed by Microsoft and since the platform maintains a minimum of two servers per technology, we ensure the security of the platform is offset against the stability of the platform.

In the event of a critical vulnerability notification from Microsoft via our Premier notification subscription we will apply any 'out of band' critical patches to all odd numbered servers immediately, test for stability and as long as no issues occur the critical patch will then be immediately applied to all even numbered servers and rebooted if required. Service availability and up time should not be affected as both servers in the relevant availability set will not be rebooted at the same time, e.g. DC1 is patched, rebooted, tested and then the same routine is done to DC2.

The Aurora IDAM solution has products patched according to individual management regimes as supplied to us by the software vendors. Where version upgrades are available, we aim to deploy point releases once a quarter, and major version updates once a year.

All patches are tested in pre-production before being applied to live. All security patches are applied to 3rd party products into pre-production upon notification from the vendor. Quality Assurance tests that form the release control for the Aurora platform are then re-run through pre-production before the patches are deployed via change control into the live platform. Service packs for the operating systems are applied in the same manner.

Operating system upgrades to the platform are maintained at the discretion of Core. Core will ensure the platform is always running a Microsoft supported operating system. Exact versions of operating system running within the platform are not made publicly available. Only customers utilising the 'Infrastructure as a Service' services of Core have access to server level details as part of their contracts.

From time to time Microsoft perform patching of the Azure fabric layer. This process involves a graceful shutdown of each Virtual Machine on the platform and a 15-minute outage window as the Virtual Machine and dependent components and GRS Storage are moved to a patched area of the Azure data centre.

All Virtual Machines deployed in an Availability Set are guaranteed to only have 1 node taken offline at a time. Because all Aurora servers, with the exception of the Sync engines, are deployed in Availability Sets the patching of the Azure fabric will not affect the availability of the Aurora service. Clients will be notified about all scheduled Azure Virtual Machine maintenance windows.

Core deploy the Aurora Platform according to the Microsoft best practice design pattern for High Availability in Azure. As the Azure fabric is responsible for the availability of the platform it is not possible to test the HA or DR features of the platform as part of a set test. Core have been running servers on Azure for over 13 years and can attest to the Microsoft vendor assertion that the Azure fabric does correctly maintain the availability of the solution during planned patching outages and during unplanned failures.

4. Service Levels

The customer, or their incumbent ITSM provider will classify the incident status at the first point of contact depending on the criteria outlined within the severity descriptions below. The customer can claim any incident severity level where it believes a higher than initially assigned severity level has been set.

Without prejudice Core will treat that incident as per the claimed priority level with immediate effect. At any point during that event, throughout the ongoing investigation, if the customer and Core are in agreement, then the priority level can be set to the correct level if it is found to be incorrect.

4.1. Aurora CIDM Service

Service Level Description	SLA	Description	Further Information
Aurora IDM Solution	99.95% Service availability		
Response and Resolution Timeframe	Severity 1 Incident **Response Time:** 15 Mins **Resolution Time:** 120 Mins	Complete loss of service affecting business critical systems	15 mins response from the point of incident being raised with Core. Membership of incident bridge upon agreement. Updates every 30 minutes throughout the lifecycle of the incident
Response and Resolution Timeframe	Severity 2 Incident **Response Time:** 30 Mins **Resolution Time:** 240 Mins	Partial loss of service that isn't business critical	30 mins response from the point of incident being raised with Core. Core to decide on severity level after initial investigation and monitoring. Hourly updates from the point of incident being raised throughout the lifecycle of the incident.
Response and Resolution Timeframe	Severity 3 Incident **Response Time:** 60 Mins **Resolution Time:** 48 Hours	Low impact, business critical systems unaffected, more than one user affected.	60 minutes response from the point of incident being raised with Core. Update provided prior to call closure for agreed sign off.
Response and Resolution Timeframe	Severity 4 Problem **Response Time:** 120 Mins **Resolution Time:** 7 days	Minimal impact to non-critical systems. Single user affected	120 mins response from the point of incident being raised with Core. Update provided prior to call closure for agreed sign off.

*** Please note SLAs on the service are for the service being available to the internet. If you choose to connect to the service via a VPN or other private network, then Core will not be responsible on the SLA of the connection as we do not control the client side equipment (firewall).*

4.2. Service Requests

Service Level Description	SLA	Description	Further Information
Response and Resolution Timeframe	Standard Service Request **Response Time:** 24 hours **Resolution Time:** 2 Working Days	A requirement that is not classified as being an incident and is not affecting service performance or availability and are not governed by DEFRA Change Control	All Service Requests will be assessed by a Core SME. If the request can be completely fulfilled within 4 hours or less of resource time, this will be delivered as part of Core's standard service offering.
Response and Resolution Timeframe	Non - Standard Service Request Incident **Response Time:** 48 hours **Resolution Time:** No SLA	A requirement that is not classified as being an incident and is not affecting service performance or availability and either IS governed by Change Control, or Non-standard in nature	All requests will be reviewed by a Core SME and the overall efforts will be calculated and costed. Details of the cost of the change will be sent to the requestor for approval. Once approved, Core will create a request to Change Management and complete the change within the agreed change window.

4.3. Incident Management

Core will actively be monitoring all of the Aurora IDAM infrastructure. This will raise automated tickets into the Core Service Desk if any component in the diagram fails to respond to ICMP or SNMP probes. Because of the high levels of redundancy in the design there are only two possible incident types. Service affecting (P1) and non-service affecting (P4). There are no partial levels of incident as all aspects of the design have automated failover to ensure continued service availability.

Service Level Description	SLA (**)	Description	Response Lifecycle
Aurora CIDM	99.95%	Availability of the solution to transform and synchronise identity attributes from or to at least one of the customers Active Directory instances (including Azure AD)	
Response and Resolution Timeframe	Priority 1 Incident **Response Time:** 15 Mins **Resolution:** 30 minutes for configuration change.	A P1 incident is defined as a full system outage. This can only occur due to complete loss of both Azure Availability zones hosting Aurora IDAM, or in the event of a full regional or global interruption to key Microsoft infrastructure and related services.	15 minutes escalation time from Core initial P1 ticket raising in their ITSM system. If issue is resolved within 15 minutes then an RCA will be sent and no MIM will be raised. Updates every 30 minutes throughout the lifecycle of the incident.
Response and Resolution Timeframe	Priority 2 Incident **Response Time:** 30 Mins **Resolution Time:** 240 Mins	Partial loss of service that isn't business critical	30 mins response from the point of incident being raised with Core. Core to decide on severity level after initial investigation and monitoring. Hourly updates from the point of incident being raised throughout the lifecycle of the incident.
Response and Resolution Timeframe	Priority 3 Incident **Response Time:** 60 Mins **Resolution Time:** 48 Hours	Low impact, business critical systems unaffected, more than one user affected.	60 mins response from the point of incident being raised with Core. Update provided prior to call closure for agreed sign off.
Response and Resolution Timeframe	Priority 4 Incident **Response Time:** 60 Mins **Resolution Time:** 7 days	No business impact.	60 minutes response from the point of incident being raised into Core ITSM via monitoring platform. Updates to Service Owner as hardware replacements delivered and installed.

4.4. Limitations

This SLA and any applicable Service Levels do not apply to any performance or availability issues:

- Due to factors outside our reasonable control (for example, natural disaster, war, acts of terrorism, riots, government action, or a network or device failure external to our data centres, including at your site or between your site and our data centre);
- That result from the use of services, hardware, or software not provided by us, including, but not limited to, issues resulting from inadequate bandwidth or related to third-party software or services;
- Caused by your use of a Service after we advised you to modify your use of the Service, if you did not modify your use as advised;
- During or with respect to preview, pre-release, beta or trial versions of a Service, feature or software (as determined by us) or due to Proof of Concept testing impacting the live service if connected through the Connectivity Hub;
- That result from your unauthorized action or lack of action when required, or from your employees, agents, contractors, or vendors, or anyone gaining access to our network by means of your passwords or equipment, or otherwise resulting from your failure to follow appropriate security practices;
- That result from your failure to adhere to any required configurations, use supported platforms, follow any policies for acceptable use, or your use of the Service in a manner inconsistent with the features and functionality of the Service (for example, attempts to perform operations that are not supported) or inconsistent with our published guidance;
- That result from faulty input, instructions, or arguments (for example, requests to access files that do not exist or to access IP addresses or resources that do not exist);
- That result from your attempts to perform operations that exceed prescribed quotas or that resulted from our throttling of suspected abusive behaviour;
- Due to your employees, agents, contractors, or vendors performing any sort of security or penetration testing of our platform without prior notice and consent from Core.

CORE

TRANSFORMING BUSINESS
THROUGH TECHNOLOGY

Core Technology Systems (UK) Limited. 9 Appold Street, London, EC2A 2AP

www.core.co.uk hello@core.co.uk +44 20 7626 0516