StonesThro

PRODUCT and SERVICE DESCRIPTION

Edge Compute Cloud Managed Services

A sovereign, low latency, distributed, neutral edge cloud that fuels business-critical data lifecycle, boosts high performance and drives revenue growth enabling game-changing possibilities.

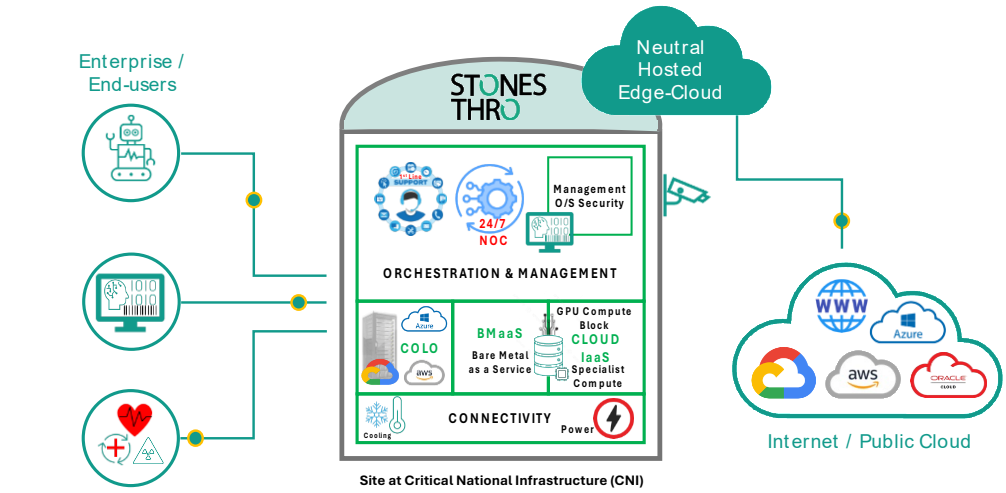

Multi-Service Orchestration & Automation

Apps securely on the right cloud

High-Performance Dynamic Networking

Connecting Clients, data centres and private clouds

Multi-Cloud Infrastructure

Interoperable & Flexible Computing Models

Revision History

Version	Date	Description
0.1	14/01/2025	First draft
0.2	24/01/2025	Additional features and use cases reviewed.
0.3	26/01/2025	Reformat the document
0.4	29/01/2025	Edits and Diagrams
1.0	31/12/2025	Release version to GCloud

Approvals

Position	Name	Signature	Date
Author:	Andy Bates Gregg Mearing		14/01/2025
Reviewed:	Mark Hua Ashley Sellar		31/01/2025

Proprietary Notice

This document is produced against the following Cornerstone output criteria

Caveats to this document

Content in this document is being evaluated with a number of potential proof of concept (PoC) customers and research conversations. The document is subject to ongoing change under a change control process as we move into PoC delivery which will inform the final productisation of the solution and services offered.

1 Table of Contents

2 Introduction

Our innovative micro edge compute service focuses on empowering businesses to innovate and optimise their operations by bringing computing power closer to where it's needed most. The solution provides the key benefits of edge computing while addressing the unique challenges faced by modern enterprises.

Key Messages

- ✓ Unparalleled **Performance**: Experience lightning-fast response times and reduced latency for your critical applications.
- ✓ Seamless **Scalability**: Easily expand your edge infrastructure as your business grows, without the complexity of traditional setups.
- ✓ Enhanced **Security** and **Compliance**: Maintain your data nearer to its source, thereby ensuring enhanced control, UK sovereignty, and adherence to data regulations.
- ✓ **Cost-Effective** Innovation: Leverage the power of edge computing without the hefty price tag of traditional infrastructure investments.
- ✓ **Future-Ready** Architecture: Stay ahead of the curve with a platform designed to support emerging technologies like 5G, IoT, and AI.

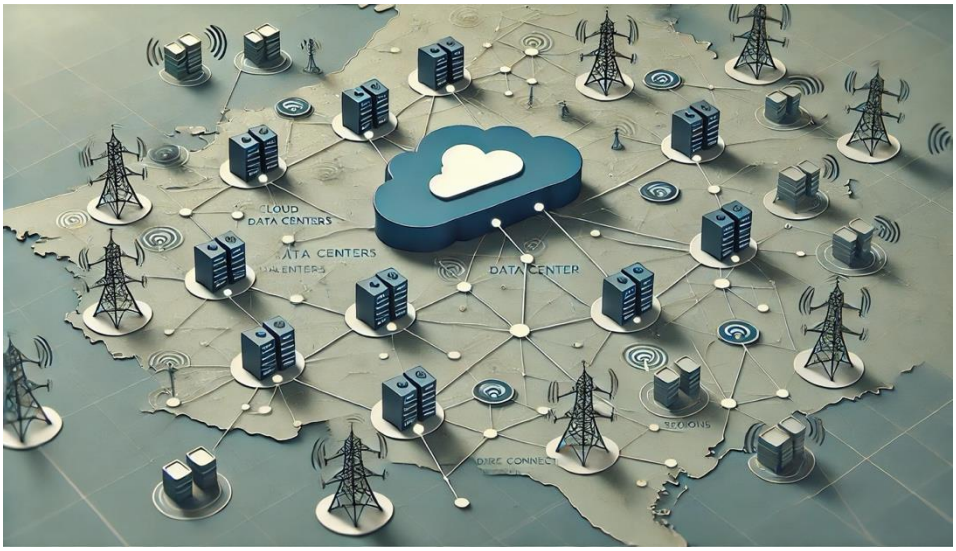

2.1 About Stonesthro

Stonesthro was founded by a team of experienced C-suite executives who have previously built, operated and deployed cloud services for SME and Enterprise customers who required highly available secure services. The services offered are delivered in conjunction with Cornerstone Telecommunications Infrastructure Limited (CTIL) a provider of critical national infrastructure services to the mobile industry across 15,000+ national locations throughout the UK. As a result of the broad geographic coverage achieved users are generally no more than 5km from a Cornerstone location.

2.2 Audience

This document is intended for:

- Technical business decision-makers evaluating cloud service options
- IT and DevOps teams responsible for the deployment and integration of production systems and services
- Security and compliance teams
- End users of the service

3 Proposition

3.1 Product / Service Name

PROXORA The Stonesthro product, however it may be rebranded by Cornerstone to their choosing for Cornerstone-initiated sales or sales to customers (such as MNO)

STORA by Cornerstone is the Cornerstone brand, with the name meaning "Stonesthro Edge" whilst remaining subtle and nondescript.

3.2 Core Product Offering

The core product offering is EmDC Virtual Data Centre (Edge micro–Data Centre) which replicates the data centre technology commonly known in the industry as a Virtual Data Centre. The difference between a commonly used Virtual Data Centre and an Edge micro Data Centre is that it is formed of much smaller infrastructure per site, and the sites number "many" as opposed to a few for Virtual Data Centres, which are usually found in Regional Data Centres.

EmDC Services are also described within this document as "The Platform".

EmDC offers:

- IaaS and PaaS services and support for third party SaaS offerings
  - CPU as a service
  - GPU as a service (for AI/ML) capability
  - Storage Services (Object, Block, File)
  - Network (Physical, Virtual, Software Defined)
  - Linux servers
  - Windows servers
  - Kubernetes as a service
  - Docker and other container support
- Colocation Services (COLO)
  Providing capacity for bespoke privately owned compute resource such as ORAN or hyperscaler local nodes

- Optional security services:
  - Virtual firewall
  - Anti-Virus/Malware and Threat Protection
  - DDOS protection
  - SOC and SIEM services

Further offerings will be available on request.

The platform can be operated as either a private, hybrid or multi-cloud service and, by design, is a multi-tenant solution.

Stonesthro provided microedge cloud nodes adjacent to mobile infrastructure with connections to hyperscaler cloud and regional/national data centres

3.3 Scalability

By offering an ability to deploy multiple microedge sites on demand gives Stonesthro the ability to scale rapidly. By deploying a blend of CPU, GPU and storage it allows us to match multiple workloads and off load high latency applications to remote nodes such that local nodes can be preserved for low latency specific applications.

3.4 Service Terms and Conditions

The offering limitations, exclusions, and/or conditions (where applicable) – are detailed in this Service Description Document for completeness. The key ones can be best summarised as:

- Security of the customer's cloud instance(s) will be a customer specification, but optional security services will be provided by Stonesthro (logical security)
- Physical security will be managed by Cornerstone
- Network connectivity to the customer's cloud instance may be subject to cyber-attack, the customer and Stonesthro will agree on defences or mitigations against these at the time of order
- The customer is responsible for any licensing of Software and operating system (as required) on top of the Stonesthro cloud services
- Stonesthro recommend that all customer production services in the Proxora environment are covered by associated disaster recovery (DR) and Backup services. These will be included in any proposed solution however they may be removed at the request of the customer.
- Stonesthro will maintain uptime to a 99.9% SLA at any one site cluster.
- The customer will be responsible for ordering backup services as appropriate
- The customer will be responsible for data processing legal requirements such as GDPR, in this case, the customer will be classified as both a processor and controller while Stonesthro is a provider of infrastructure.
- Further details are provided in our Terms and Conditions document which take legal precedence over this document.

4 Pricing Model and Principles

Minimum usage spend and committed services (Reserved Resources) will be invoiced in advance on a monthly basis. Consumption based and overage (resources consumed over committed or minimum usage levels will be invoiced monthly in arrears. Discounts are offered for Reserved Resources (Commit), Minimum Usage and contract term of more than 12 months.

Pricing for compute resources will be charged according to:

- CPU size
- GPU size
- RAM
- Storage
- Response latency (a roadmap option), see roadmap section later.

Discounts are available for long-term commitments to workloads after several months.

In the future, a premium may be applicable for low latency response times to further differentiate the quality of service (QoS).

Discounts are available for "batch processing" allowing customers to use resources during periods where excess unused capacity is available. These will be discussed and agree with the customer during the proposal or ordering process depending on the specific requirements and conditions required to utilise such services.

Prices may be increased at any time following the expiry of the initial contract term and will be based on the greater of RPI or 3%. Additionally, the price may be increased, subject to written notice of 30 days or more to the customer, where costs incurred to deliver the service are increased outside of our control including, but are not limited to, electricity, regulatory changes, third party provided services including software licensing.

Variable pricing, such as licensing and power, shall be subject to price variance against a benchmarking regime.

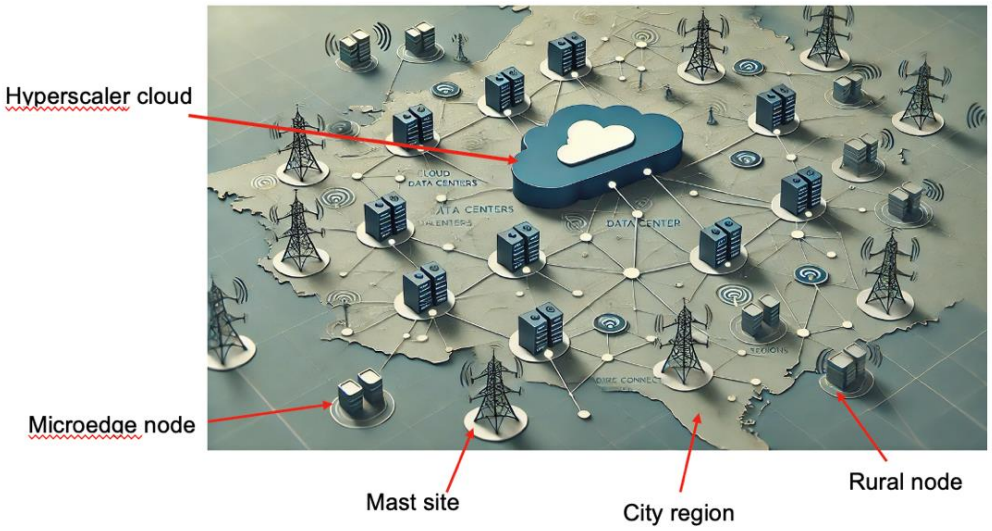

4.1 Backup Principles

The Platform may be used as a backup for customers' other IT infrastructure. The backup solutions offered are:

- Single site node
- Dual node in local proximity (meshed cluster(s)/microregion(s))
- Dual node in regional proximity (geo-resiliency)
- Triple node (one local and one national)

Oxfordshire for example would be considered a Microregion.

In all cases, Stonesthro will encourage customers to take some form of backup either by Stonesthro or by other parties to assure resilience against single node failures.

All services shall be backed up by Stonesthro unless backup removal is explicitly requested by the customer. Service pricing discounts are available in these cases.

4.2 Portal

Customers will order and manage the service via the Stonesthro portal.

We also have an API which complies with the EC2 standard. This allows developers familiar with AWS cloud to be able to utilise this interface without re-tooling.

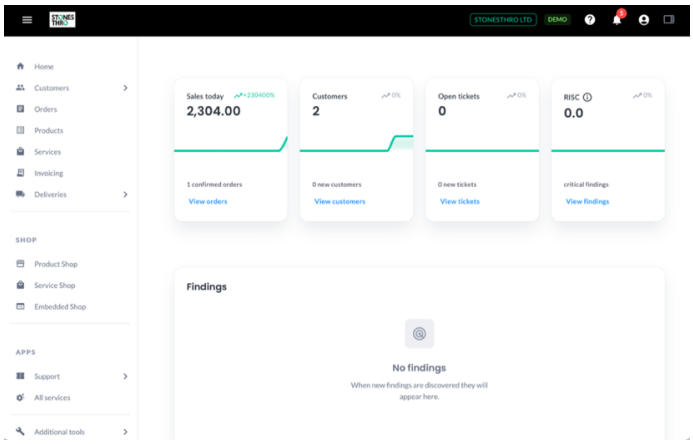

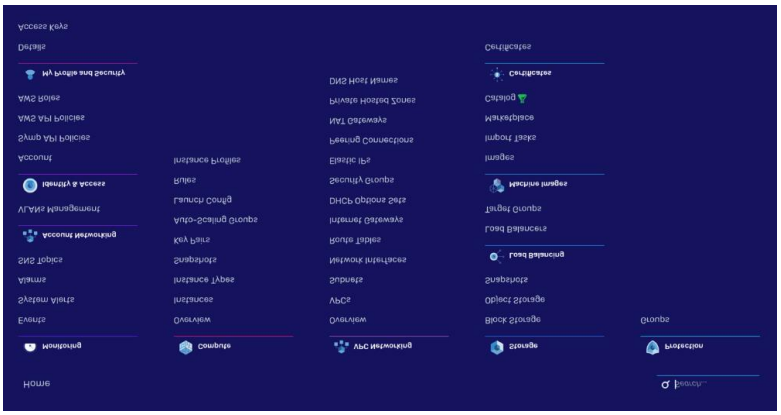

4.3 Service Interface

Service outages and incidents will be reported by the portal to customers based on the impact categorisation tabulated below.

Severity	Ack Time	Priority	Impact Level	Response Time	Investigation and Remedy	Update Frequency
High	15 mins	P1	Total service affecting	15 mins	2 hrs	1 hr
Medium	1 hr	P2	Potential for business impact or partial service impact	1 hr until agree to pass to next working day	4 hrs	2 hrs
Low	4 hrs	P3	Low impact or ancillary service	4 hrs	8 hrs	Next working day
Info	Next working day	P4	Information or Ad-Hoc	Next working day	N/A	Following working day

Other roadmap functionalities may include:

- Ticket integration to customer ticketing systems such as JIRA will be deployed as roadmap functionalities.
- Live chat is also available; and
- Helpdesk (Phone) Support.

5 Features and Benefits

This section is broken into two sections, conventional cloud use cases and edge cloud use cases. Stonesthro delivers conventional cloud as well as the specific edge cloud capabilities.

5.1 Cloud Benefits

Cloud offers:

- Consumption-based charging
- Avoidance of regular capex hardware spend
- Flexibility to deploy and change compute resources (changing either capacity or type of resource).

5.2 Edge Cloud USPs

Edge cloud can perform all the functions of conventional cloud, in addition to some physical and logical aspects, that can only be effectively delivered by compute resource located close to the user and workloads they require:

- Ultra-low latency and high throughput for time-sensitive applications
- High availability, resilient, distributed architecture design
- Hyper proximity allowing compute resource(s) to be located within the appropriate distance of the customers facility
- High ESG credentials such as a PUE as low as 1.1 and the ability to offer waste power to the local community

- Agile deployment of new nodes dependant on geographical requirements
- Provide extension services for regional and global cloud services, agnostically supporting multiple and hybrid cloud deployments
- Broad distribution of federated compute resources providing greater resilience to physical compromise
- Services delivered exclusively from UK based facilities, guaranteeing data sovereignty
- Edge processing offload to optimise and reduce required network resources and associated costs
- Full encryption of data in motion and at rest with further, higher grade security options available

---

6 Security and Regulations (Infosec)

All cloud services are delivered from facilities physically located within the UK.

All facilities are national infrastructure premises provided from:

- Sites within a mobile phone estate of 15,000+ locations
- Tier-3 data centres from approved suppliers.

Operation of the cloud will be by UK-based staff subject to the relevant vetting and security clearances.

Stonesthro is Cyber Essentials certified and will be ISO27001, ISO9000 certified and aligned to ISO27017 and ITIL standards. The scope of our ISO certification shall encompass the core cloud components, network, management and staff including third parties.

Our material subcontractors and third parties are ISO27001 certified.

Cornerstone sites can, where necessary, meet the physical security requirements, including remote access control and monitoring, to the level required to host critical national infrastructure (CNI) designated services.

Stonesthro's Cloud infrastructure will be actively monitored for both logical and physical intrusion.

Bespoke security clearances will be available on demand for highly classified operations.

Encryption or data in transit and at rest will be included and can be elevated to customer bespoke specifications on request.

500m pa]

7 FAQ – Frequently asked questions

Will traffic ever leave the UK?

No Stonesthro is an entirely UK-based infrastructure with a captive network and encryption technology.

How many sites could cloud be deployed within?

We have access to 15,000 sites although 10,000 seem suitable for cloud deployment

What is the PUE of your services (power unit efficiency)

At launch we will be PUE of 1.02 which is beyond industry best practice. In reality many hyperscale data centres claim a PUE of 1.2 but do not achieve it until 90% full. Micro edge data centres are in essence built at 100% capacity per site so achieve the desired PUE much sooner than comparable hyperscalers approaches.

What is the difference between a cluster and a swarm?

A cluster is where there are two or more microedge nodes provide backup and resilience to each other.
A site cluster is where two or more sites contain nodes that are joined as a single entity providing resilience against both single node and single site loss.
A swarm is a number of nodes normally in a region, an example could be 5 nodes in Oxfordshire. This provides high availability and low latency and increased resource flexibility. A swarm is effectively a large cluster.

How do I know my data is secure?

Stonesthro team have a vast range of experience in the cloud and cyber security industry and have worked in line with industry best practice. However rather than rest on the current status quo Stonesthro have also embraced innovate technology such as data sharding and quantum proof encryption principals as well as partnering with a CNI infrastructure partner to produce a holistically secure solution.

How does microedge cloud increase my ESG credentials?

By having cloud closer to users and distributed throughout the UK makes a better use of the UK national grid infrastructure and by lowering PUE reduces the overall net power consumption.

8 Technical Detail

8.1 Compute Instances (VMs)

Virtual Machines of any type (built-in or user defined) can be created based on an OS image with the use of a cloud-optimised hypervisor (Enhanced KVM). This service has AWS EC2 compatible API that offers standard AWS API for application integration.

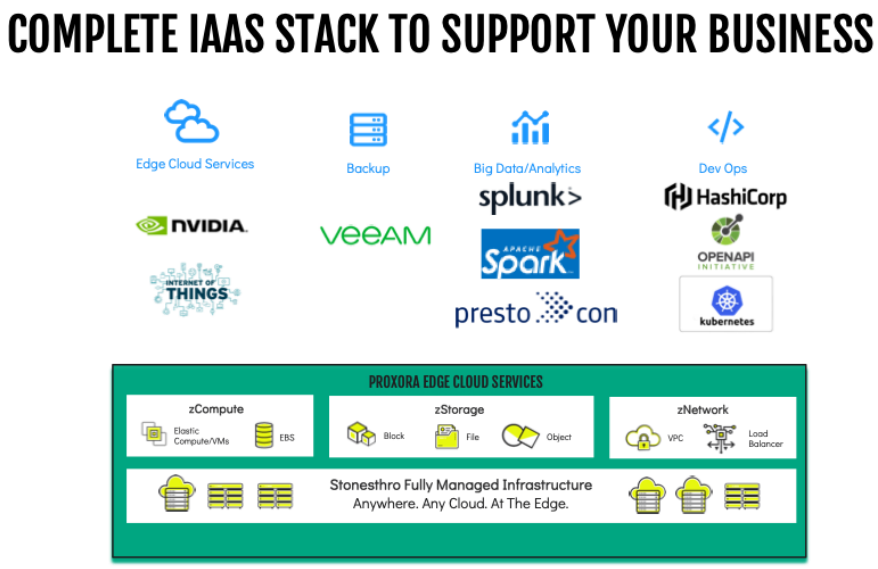

8.2 Snapshots

A snapshot is a copy of the volumes of a virtual machine at a given point in time. Snapshots provide a change log for the volume and are used to restore a VM to a particular point in time when a failure or system error occurs. You can create snapshots of the VMs that are in the system and use these snapshots to restore any VM if needed.

8.3 Key Pairs (Security Key Settings)

Public–key cryptography is used to encrypt and decrypt login information for Linux instances. (Instances created from a Windows image are accessed by a password instead.) A public key is used to encrypt a piece of data, and then the recipient uses the private key to decrypt the data. These two keys are known as a key pair. Key pairs enable you to securely access your instances using a private key instead of a password.

8.4 Auto Scaling

The system monitors applications and automatically adjusts allocated resources to meet evolving requirements. It simplifies scaling operations to a minimal e ort based on a simple and powerful interface for planning and defining launch configurations and

scaling policies. You define the policy, and the system ensures performance and availability while keeping resource utilisation optimised to reduce costs.

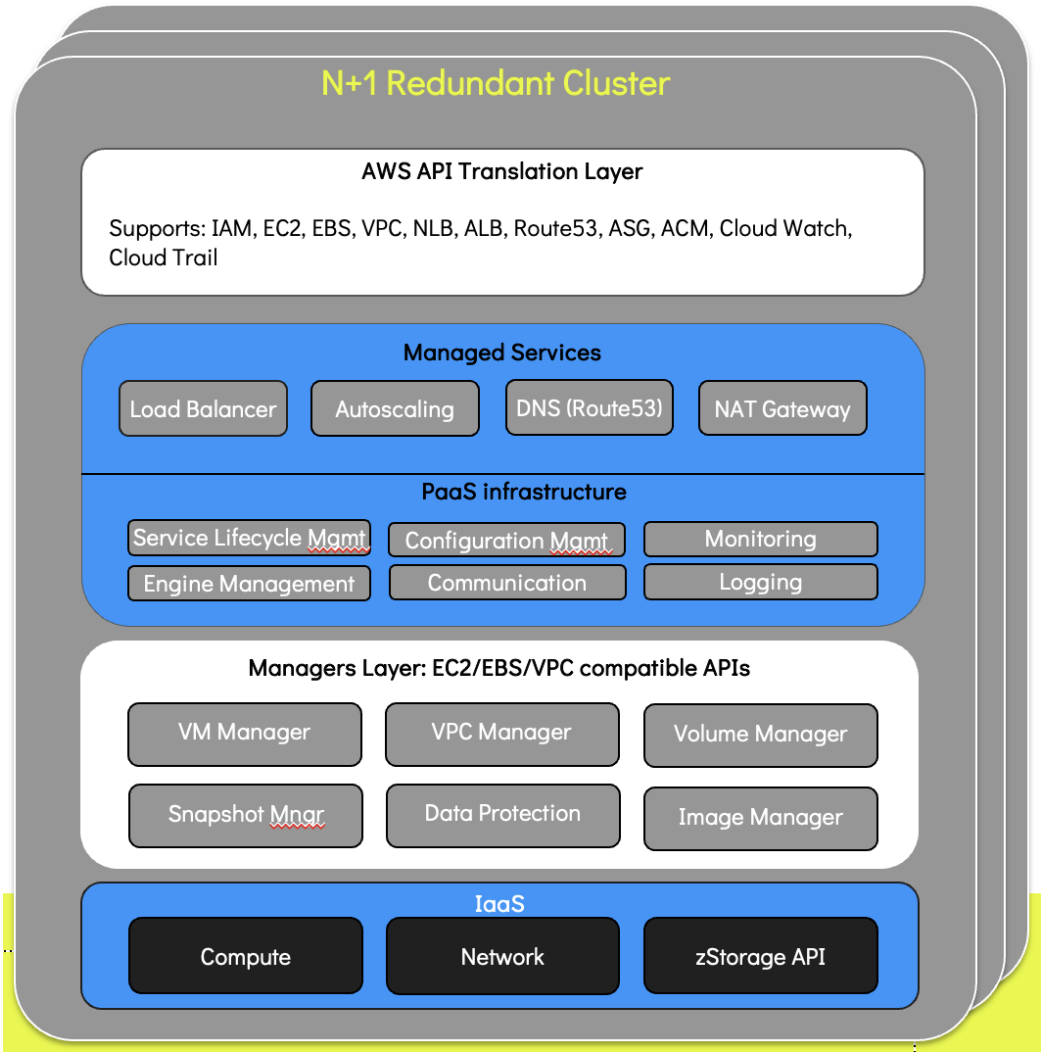

9 Networking

9.1 Virtual Public Cloud (VPC)

The Virtual Public Cloud (VPC) is your private data centre within the cloud infrastructure. You get to select the network addresses that will be used throughout your infrastructure. Since this is your network, you can slice it up any way you prefer. The VPC is a networking resource with a logical router at its core. When you create a VPC, you specify a CIDR block. All subnets created in the VPC will be carved out from this CIDR block, and the router will ensure IP connectivity between them.

9.2 Subnets

Subnets are used for the configuration of networking within a VPC. Compared to using one big network, using subnets in VPC allows segregation between private and public-facing networks, definition of availability zones, and other similar advantages.

9.3 Network Interfaces

A virtual network interface (VIF) is an abstract virtualised representation of a computer network interface. Each virtual machine (VM) may have one or more virtual network interfaces that act as virtual network interface cards (NICs).

9.4 Route Tables

The basic tenet of networking is that everything inside the subnet stays inside the subnet – and if one wants to go outside the subnet, it will need to go through the default gateway through which traffic is routed to an external destination network. The route tables are associated with each subnet to allow the flow of traffic according to the VPC policies and configured options.

9.5 Internet Gateway (IGW)

Connection to the outside world is via the Internet Gateway. Having instances running in a cloud is great, however, if the outside world cannot get to them, then it might be useless. Without an IGW, manageability would be very challenging – if not completely impossible.

9.6 Elastic IP (EIP)

Elastic IPs are used to expose instances outside of the Stonesthro compute cloud. An EIP will be used in the network address operation (NAT) of all traffic to/from the virtual network interface with which it is associated.

9.7 VPC Peering (in the same cluster)

VPC peering lets you create direct IP connectivity between any two VPCs. Direct connectivity between VPCs means that servers in a VPC can be reached from the

other VPC without the need for elastic IPs or traffic flowing through external routing. VPC peering is simply L3 connectivity realised using routing tables and IP connectivity.

9.8 NAT Gateway (NGW)

There will be cases where one does not want instances to be exposed to the outside world, nor want them to have public IP addresses. However, these instances may still need access to the outside world to get updates or to send information. By using an NGW, outbound access to the internet can be allowed, while limiting inbound access, thus providing an additional layer of abstraction and protection for your workloads.

9.9 DNS

The DNS Service provides an easy-to-manage DNS solution for your private network. It offers high availability and a cost-effective way to connect applications and services and make them available to users. The DNS service is seamlessly integrated with the resources and services of the cloud. Domain names to compute instances and other services can then be map.

9.10 Security Groups

Security groups are essentially whitelists applied to the virtual network interfaces to control the inbound and outbound traffic. Traffic that does not match any rule in the security group will be discarded.

For each security group, you include one set of rules that controls the inbound traffic to the instances and a separate set of rules that controls the outbound traffic from the instances.

---

10 Storage

10.1 Block Storage (EBS)

EBS (Elastic Block Store) is native block storage, designed to be extremely scalable, agile and flexible while being fully compatible with the requirements of cloud computing. The Block Storage service offers comprehensive block storage capabilities, enabling reduced latency and overhead while increasing throughput. The service is based on Stonesthro VPSA block services and is fully compatible with AWS EBS API for application integration.

Different capabilities are available, based on Volume types.

10.2 Volume Types

Stonesthro Cloud Platform simplifies storage management with volume types that provide users with a range of options to meet their storage requirements, balancing factors such as performance, cost, and specific workload demands.

By abstracting the underlying storage infrastructure, volume types simplify storage management and allow users to focus on selecting the appropriate type for their Stonesthro Cloud Platform needs.

Different volume types offer varying levels of input/output operations per second (IOPS) and throughput, which determines the storage performance. Higher performance volumes are typically associated with faster data transfer rates and lower latency, enabling applications to handle demanding workloads or perform intensive operations.

10.3 File Shares

NAS service is available to provide scalable les shares via NFS or SMB protocols. The service is based upon the Stonesthro VPSA NAS offering allowing access to all associated VM instances. (This service is currently provisioned via the Stonesthro VPSA interface and is incompatible with AWS EFS API).

10.4 Object Storage

Both private and public object storage services are available on the Stonesthro IaaS platform. Private Object Storage can be created where one can have full control, or just an account on a Stonesthro regional public service. Both can be accessed via AWS S3 API or Swift interface.

A private Object Storage is directly connected to the VPC, while the public Object Storage is accessible via a public IP. (This service is currently provisioned via the Stonesthro Object Storage interface).

---

11 Load Balancing

11.1 Load Balancers

The Managed Load Balancer service offers the ability to spin up, customise and scale load balancers to support fault tolerance and ensure high availability and application scalability over time. To simplify operations even further, this service supports AWS ELB APIs. It routes traffic according to application or network considerations, provides the required amount of load balancing capacity needed, and distributes it to meet high availability and network performance requirements.

11.2 Target Groups

A target group is a group of instances to which a load balancer directs application traffic. The instances in this group collectively do the processing work that the application requires.

---

12 Machine Images

12.1 Images

The system comes with several VM images ready to use. In addition, one can create images as needed. Two image formats are supported - VMware OVA image and KVM-compatible image in RAW or QCOW2 formats. The latter means that most OpenStack KVM images can be uploaded easily. All the above are converted into the RAW format and preserved in the image repository in the cluster.

12.2 Images Marketplace

Machine images of the different OSs can be downloaded from the Images Marketplace and be used to create Instances. For OSs that require activation (and vendor fees), you will need to get the activation codes directly from the vendor.

---

13 Monitoring and Logging

13.1 Events

Stonesthro compute cloud provides information about various events in your system. Events are any actual changes or attempts at changes in the state of the services, whether initiated manually or automatically by the system. This service provides an events viewer with the ability to filter by date, severity, time, and other relevant parameters.

The system offers visualised metrics, logs and events of all resources, including CPU utilisation, memory, storage and network usage. Insights are available on a broad system view, as well as on a per-instance and per-volume view, to quickly and effectively resolve issues by analysing diagnostics to understand the root cause.

You can configure the system events to be sent to a remote syslogs/logstash server located at some external endpoint.

13.2 Alerts

The cloud's automated monitoring helps focus on proactively solving problems rather than wasting time to identify them. You can set alarms on all relevant metrics and reduce the response time for facilitating automation. A set of metrics-based alarms with multiple conditions to notify threshold crossing can easily be customised to provide real-time alerts.

---

14 Certificates

To secure a cloud domain name, an SSL Certificate from a trusted SSL Certificate Authority needs to be procured. Once the certificate is available from CA, the service then allows the certificate to be applied to the system.

---

15 Protection

15.1 Group Policies

The Protection Group essentially consists of a schedule for automatic local backup of a project. It is a backup policy that defines the backup window, recurring schedule, and retention period.

In addition to the member services above, the following services are available for Tenant Admins.

---

16 Identity and Access Management

16.1 Account

Tenant Admins have the added ability to monitor and modify the account state, create and manage projects within the account, and member user's access rights.

16.2 Access Keys

Tenant Admins can manage account-specific credentials that authenticate users before they can access the system. For example, to access the GUI or Management console, users sign in with a username and password. Similarly, for programmatic access, users provide an access key and a secret key. As the API's are AWS compatible, the authentication mechanism is also identical

16.3 Roles

Roles are policy-based tokens with temporary credentials allowing a user temporary access to AWS services and actions which the user is normally not permitted to access. These users may be from different projects or even different accounts. These roles can also be embedded into specific instances, allowing these instances access to the necessary actions.

16.4 API Policies

Usage of all compatible services and actions are governed by their corresponding managed policies. These policies can be assigned per project to Users, Groups of users and STS Roles. System usage is governed by Stonesthro Cloud Services policies together with the Stonesthro roles.

---

17 Edge Cloud Services Detail

17.1 Edge Compute Services

17.1.1 Overview

Our Edge Compute Services bring processing power closer to where it's needed, enabling faster response times and reduced latency for critical applications.

17.1.2 Key Features

- High-Performance Computing: Deliver low-latency application requirements with our distributed network.
- Scalable Resources: Access high-density, high-performance compute resources on-demand.
- Flexible Instance Types: Choose from a range of pre-selected instance sizes optimized for various workloads.
- Secure by Design: ISO 27001, 27017, and 27018 certified or aligned with built-in security features.

17.1.3 Benefits

- Reduced latency for time-sensitive applications
- Improved application performance
- Cost-effective scaling of compute resources
- Enhanced data security and compliance

17.2 Edge Storage Services

17.2.1 Overview

Our Edge Storage Services provide fast, reliable, and secure data storage solutions at the edge of the network, ensuring quick access and processing of data where it's generated.

17.2.2 Key Features

- High-Performance SSD Storage: Utilize our SSD storage platform for superior performance.
- Flexible Storage Tiers: Choose between Performance and Capacity tiers based on your needs.
- Data Protection: Benefit from integrated backup and disaster recovery options.
- Encrypted Data at Rest: Ensure the security of your stored data with encryption.

17.2.3 Benefits

- Reduced data transfer costs
- Improved data access speeds
- Enhanced data security and compliance
- Flexible storage options to match various workloads

17.3 Cloud Orchestration Services

17.3.1 Overview

Our Cloud Orchestration Services provide seamless management and integration of edge and cloud resources, enabling efficient hybrid and multi-cloud deployments.

17.3.2 Key Features

- Hybrid Cloud Integration: Seamlessly connect edge resources with major public cloud providers.
- Automated Deployments: Enable rapid deployment of workloads at a national scale.
- Cloud Cost Optimisation: Manage and optimize cloud expenses effectively.

17.3.3 Benefits

- Unified management of edge and cloud resources
- Improved resource utilization and cost control
- Faster deployment and scaling of applications
- Enhanced flexibility in workload placement

17.4 AI Services at the Edge

17.4.1 Overview

Our AI Services at the Edge enable businesses to leverage artificial intelligence and machine learning capabilities closer to the data source, facilitating real-time insights and decision-making.

17.4.2 Key Features

- Edge-Optimised AI Infrastructure: Purpose-built hardware for AI workloads at the edge.
- Low-Latency Inference: Enable real-time AI applications with reduced latency.
- Distributed Machine Learning: Run machine learning models across edge locations for improved performance.
- AI-Ready Data Processing: Prepare and process data for AI applications at the edge, with data encrypted and held locally to ensure privacy and protection.

17.4.3 Benefits

- Faster AI-driven insights and decision-making
- Reduced data transfer costs for AI workloads
- Improved privacy and data sovereignty for AI applications
- Support for emerging AI use cases in IoT and 5G environments

17.5 Service Assurance

All our Edge Cloud Services are backed by:

- 24/7 Expert Support: Round-the-clock assistance from our UK-based operations team.
- Geo-Resilient Infrastructure: Services available across multiple UK regions for enhanced reliability and disaster recovery.
- Consumption-Based Billing: Pay only for the resources you use, thus optimising the IT spend.
- Compliance and Security: Meet strict data protection and security standards with our certified services.

We will provide level 1 and level 2 support with escalation to c-suite in the event of all priority 1 incidents

17.6 Maintenance and Updates

It is assumed that every node will require 1 hour of down time per month. This may not be required on every month. Customers will be informed prior to service disruption. As far as its practical updates will be at one node per time to allow resilience within customers designs to maintain service. Information will be available via the portal and email (it is suggested that functional team emails are used rather than personal emails).

17.7 Emergency maintenance:

Where a critical security event occurs Stonesthro will invoke an emergency patch regime which may be service affecting but will be undertaken to preserve the security of the service and customers using the service. In these cases maintenance may occur without prior notification.

17.8 Security patching and hardening.

Patching and hardening will be undertaken during the maintenance windows above.

Patching will be in line with industry best practice ensuring customers maintain compliance with NCSC cyber essentials and ISO27001.

18 Premium Support and Professional Services

Customers can self-provision using the management portal however we can offer the following services, at an additional cost. These will be quoted on request and priced depending on the size and complexity of the services required.

- Sign-up and Account Creation
- Resource Allocation and Configuration
- Migration and Integration
- Data Migration Services
- API and SDK Support
- CI/CD Pipeline Integration

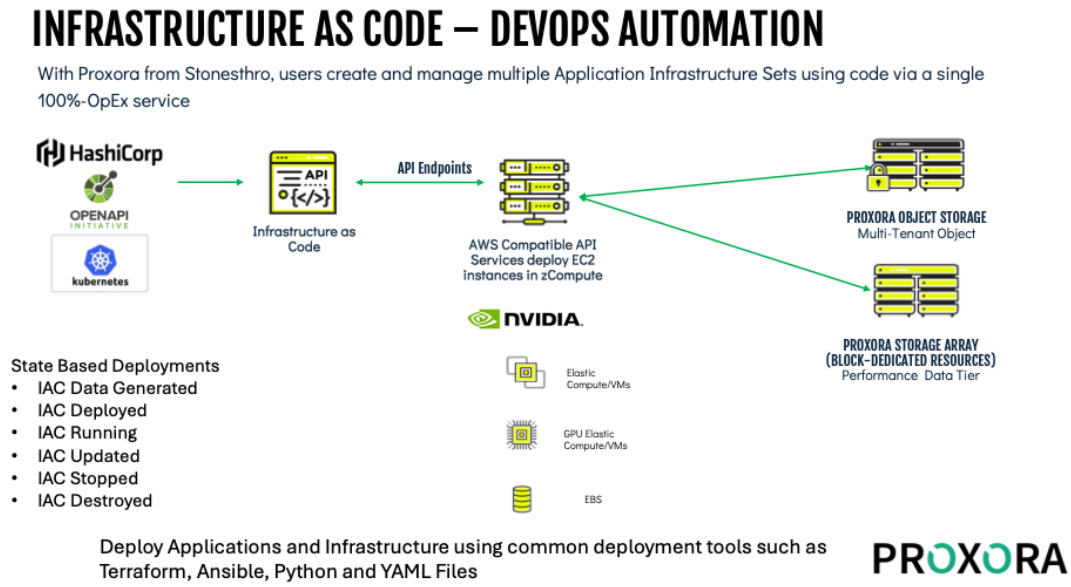

19 Roadmap

All our services will evolve in line with market expectations and be ahead of predicted demand, this section covers the current roadmap informed by current conversations and market research.

19.1 Latency-Based Pricing Principles

The following are latency pricing principles which will be tested in POC and market research.

Service Level	Latency	Pricing Principles
Platinum	<5ms latency	+30% premium
Gold	<10ms latency	+20% premium
Silver	<20ms latency	+10% premium
Bronze	<50ms latency	Cloud market price

19.2 Compute Heat Recovery

Local heating is a roadmap option informed by proof of concept.
NPU Neural processor units will be a roadmap item once GPU have been established in the microedge domain. A NPU is a processor designed for low power low form factor AI originally designed for AI in mobiles.

19.3 API Ticket Integration

Ticketing support by API such as JIRA will be evaluated post proof of concept.

19.4 GPUaaS

The ability to deliver GPU as a Service at the Edge is another thought leadership from Stonesthro. By enabling a combination of different GPU models to handle different workload tasks, be that AI/ML with NVIDIA vCS drivers and GPU's, Graphics intensive workloads for Remote workstations or providing a remote desktop Office environment. Stonesthro is going to enable this with its deployment of Edge Storage and Compute.

---

20 Glossary of Terms

A

- **API (Application Programming Interface)**

A set of protocols and tools that allow different software applications to communicate with each other. Used to integrate services within Stonesthro's cloud environment.

- **Auto Scaling**

A cloud computing feature that dynamically adjusts computing resources based on workload demand to optimize performance and cost efficiency.

- **Availability Zones**

Physically separated data center locations designed to enhance redundancy and prevent service disruptions.

B

- **Backup and Restore**

A cloud-based disaster recovery mechanism ensuring data integrity by maintaining copies of data at different locations.

- **Bare Metal as a Service (BMaaS)**

A cloud service model providing direct access to physical servers without a virtualization layer.

C

- **CDN (Content Delivery Network)**

A geographically distributed network of servers that work together to deliver digital content quickly to users by caching it closer to their location.

- **Cloud Computing**

The delivery of computing services—including storage, processing power, and networking—over the internet.

- **Colocation Services**

Hosting of privately owned servers and networking equipment in a third-party data centre.

- **Compute Instances (VMs)**

Virtualized computing environments that provide on-demand processing power for cloud-based applications.

- **Cyber Honey Pot**

A security mechanism designed to attract cyber attackers to a decoy system to analyse their behaviour and enhance cybersecurity measures.

D

- **Data Sharding**

A database partitioning technique that splits large datasets into smaller, manageable pieces distributed across multiple storage locations to improve security and performance.

- **DNS (Domain Name System)**

A hierarchical naming system that translates domain names into IP addresses to facilitate internet communication.

- **DDoS Protection (Distributed Denial-of-Service Protection)**

Security measures designed to prevent and mitigate attacks where multiple compromised systems target a single service.

E

- **Edge Computing**

A distributed computing model that processes data closer to its source rather than relying on centralized cloud data centres, reducing latency and bandwidth usage.

- **Edge micro-Data Centre (EmDC)**

A small-scale data centre located close to end-users or devices to enable low-latency data processing.

- **Elastic IP (EIP)**

A static, public IP address designed for cloud environments to provide consistent connectivity for applications.

F

- **Firewall (Virtual Firewall)**

A network security system that monitors and controls incoming and outgoing traffic based on predefined security rules.

G

- **GPU Compute (Graphics Processing Unit Compute)**

The use of GPUs to accelerate computing tasks, particularly for AI, machine learning, and high-performance computing applications.

- **GPU as a Service (GPUaaS)**

A cloud-based service providing access to GPU resources for intensive computing tasks like AI and deep learning.

H

- **Hybrid Cloud Integration**

The combination of on-premises, private cloud, and third-party public cloud services to create a unified computing environment.

I

- **IaaS (Infrastructure as a Service)**

A cloud computing model that provides virtualized computing resources over the internet, including servers, storage, and networking.

- **Internet Gateway (IGW)**

A network component that enables internet access for cloud-based applications.

- **ISO 27001**

An international standard for managing information security.

L

- **Latency-Based Pricing**

A pricing model that charges based on the response time of cloud services, offering different price tiers for different latency levels.

• **Load Balancing**

The process of distributing network traffic across multiple servers to improve availability and reliability.

M

• **Machine Learning (ML)**

A subset of artificial intelligence that enables systems to learn from data and make predictions or decisions.

• **Managed Load Balancer**

A service that automatically distributes incoming traffic across multiple instances to enhance performance and reliability.

N

• **N+2, N+3, N+4 Hyper Resilience**

A redundancy model ensuring high availability by maintaining additional backup systems to prevent failures.

• **NAT Gateway (Network Address Translation Gateway)**

A network function that enables instances in a private cloud network to connect to external services while keeping their IP addresses hidden.

O

• **Object Storage**

A cloud storage model that stores data as objects rather than in a hierarchical file system, optimized for scalability and retrieval efficiency.

• **Orchestration and Automation**

The automated management and coordination of cloud computing resources and workflows.

P

• **Private Cloud**

A cloud computing environment dedicated to a single organization, offering enhanced security and control.

• **Proxora**

The brand name of Stonesthro's edge computing platform.

R

• **Route Tables**

Networking configurations that determine the paths that network traffic follows in a cloud environment.

S

• **Security Groups**

Virtual firewalls that define inbound and outbound network traffic rules for cloud instances.

• **Snapshots**

Point-in-time copies of virtual machines or storage volumes used for backup and recovery.

• **SOC and SIEM (Security Operations Centre and Security Information and Event Management)**

A framework for detecting, analysing, and responding to cybersecurity threats.

• **Software-Defined Networking (SDN)**

A cloud-based approach to managing network infrastructure through software-based controllers.

• **Swarm Computing**

A network model where multiple edge computing nodes work together to provide increased performance and reliability.

T

• **Tenant Admins**

Users with administrative privileges for managing cloud resources in a multi-tenant environment.

V

• **Virtual Data Centre (VDC)**

A cloud-based infrastructure that replicates the capabilities of a traditional data centre, offering scalable computing, storage, and networking resources.

- **Virtual Public Cloud (VPC)**

A cloud networking framework that provides an isolated, private environment within a public cloud.

W

- **Workload Scaling**

The ability to dynamically adjust computing resources based on demand to optimize performance and cost.