Public Cloud

4D Public Cloud Hosting: IaaS

Technical Overview

Table of Contents

Public Cloud

Service Overview

Service Description

4D Public Cloud Hosting is based within our highly secure, UK data centres and supported by our UK support and operations teams.

4D Public Cloud VMs are built on a scalable, secure and open source infrastructure using the KVM hypervisor and isolated storage for each VM. The platform is certified to ISO 27001 along with all facilities and operations by 4D.

The public cloud platform uses shared infrastructure (hypervisors, network and storage) but with secure separation between each VM including our four layer CIM (Customer Isolation Module) to ensure that no one VM can access any other VM traffic unless explicitly allowed.

Multi-tiered storage is available for all VMs with options on 7.2k SAS, 15k SAS and SSD depending on your IO requirements allowing you to use appropriate storage for your application requirements.

One of the main benefits of the 4D cloud platforms is that VMs can be interconnected between the public / private clouds as well as back into dedicated colocation or existing networks allowing for secure data transfer and flexible deployment to meet security requirements

Features

- Online self-service portal allows you to easily deploy VMs with a few clicks and provides full console access via the web portal
- Full API integration
- Boot from ISO images or create a backup of a provisioned VM and convert to a template for rapid re-deployment
- Full isolation of VMs, network and storage. Our CIM (customer isolation module) prevents network spoofing and traffic sniffing while all storage is provisioned in dedicated LUNs
- Average 7-minute provisioning times for Linux and Windows VMs with a range of preconfigured, sysprep'd templates available for VMs covering a wide variety of operating systems

- Ability to create custom ‘recipes’ that can be applied to VMs once built allowing for the installation of custom packages, applications and security configurations
- Application Server deployments delivering ‘SaaS’ services for a wide range of web applications from content management systems (Wordpress, Drupal, Joomla) through to e-commerce packages and social networks
- Integrated content delivery network (CDN) using the ‘4D Accelerator’ instantly push your VM out to our global CDN as well as having your images and HTML pages optimised on the fly
- Full root or administrator access to VMs along with console access through our online portal
- Integrated rules based stateful firewall for each VM, fully configured through the online portal. Allow only the traffic you want to the VMs including between VMs on private networks
- 99.999% network uptime SLA and 24x7x365 monitoring of the platform from our UK based support centre
- Eco-friendly: Our platforms are hosted within some of the most energy efficient data centres in the UK with an average PUE of 1.14
- Instant backups, take manual snapshots or schedule snapshots of all your VMs. Snapshots can also be converted into new templates for quick re-deployment of similar configurations
- ISO 27001:2013 certified data centre, operations and cloud platform.

Public cloud instances – example costs

Small Instance	
Set-up	£0.00
Monthly	£33.00
OS	Linux (Ubuntu, Debian, RHEL, CentOS, Gentoo)
vCPU	2
RAM	2 GB
OS disk: 15k SAS HDD	60 GB
Data transfer	500 GB
Backup space	100 GB

Medium Instance	
Set-up	£0.00
Monthly	£53.00
OS	Linux (Ubuntu, Debian, RHEL, CentOS, Gentoo)
vCPU	4
RAM	4 GB
OS disk: 15k SAS HDD	60 GB
Data transfer	1000 GB
Backup space	200 GB

Large Instance	
Set-up	£0.00
Monthly	£83.00
OS	Linux (Ubuntu, Debian, RHEL, CentOS, Gentoo)
vCPU	6
RAM	8 GB
OS disk: 15k SAS HDD	60 GB
Data transfer	2000 GB
Backup space	500 GB

SPLA based licensing for Microsoft software is also available (Windows Server 2008/1012, MS SQL 2008/2012 Standard or Enterprise, and Exchange)
Additional disk space is available on all VMs. Example: 100 GB SSD = £10/month, 100 GB SAS 15k = £5/month, 100 GB SAS 7.2k = £3/month

The Platform

Overview

The 4D Cloud platform has been built from the ground up to provide the best mix of flexibility, resilience, security and performance. Test it out yourself with a free 14-day trial.

Platform Diagram

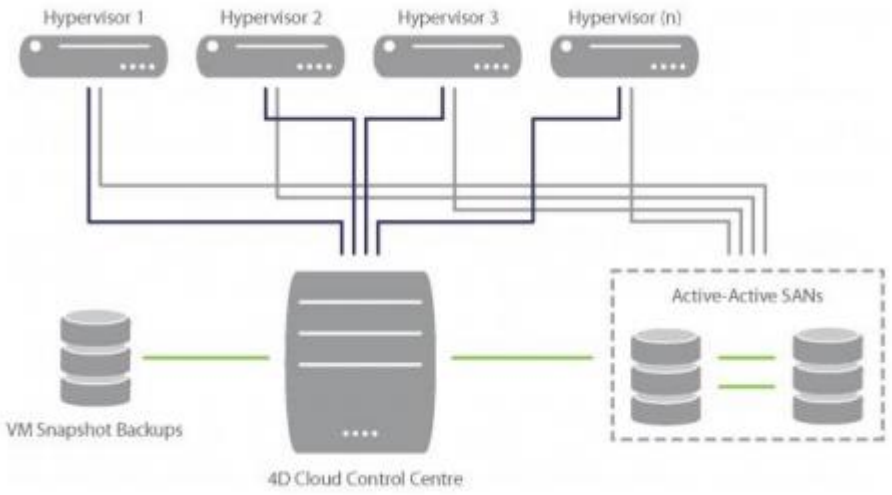

Hypervisors

Each node in the 4D Cloud has 16 dedicated cores (supporting a maximum of 64 vCPUs at 1:4 contention), 192 GB ECC RAM (supporting a maximum of 192 GB of virtual memory at 1:1 contention) and runs the 'KVM hypervisor'. After a careful review of all hypervisors available, KVM was picked for the following reasons:

- KVM supports Linux, Windows, Solaris and BSD as "to-boot". There is no need for additional drivers improving the stability of the virtual server operating system
- KVM also brings a host of performance benefits, especially for Windows, such as less I/O latency to the underlying hardware and removing the proxy between hypervisor operating system and virtual machine operating system. This means the round-trip time to the underlying hardware from the virtual server is significantly reduced

Network

Over the last seven years we've built a Cisco ASR based network for our data centres, lit our own fibre, designed networks for our clients and provided high speed connectivity to their offices.

All this experience was brought into designing the network for the 4D Cloud platform to ensure it is scalable and resilient.

The 4D Cloud platform has segregated networks for storage, internal traffic and external traffic with dual paths for each through separate Cisco switch stacks. This ensures that the platform stays up if a switch dies or if we have to perform maintenance.

The Internet connectivity into the platform is protected by a pair of WatchGuard firewalls providing intrusion prevention and detection for the whole platform. Traffic is then routed by the 60 Gbps 4D core network through a multihomed feed of Level3, NTT, Telecom Italia, Lonap and LINX routes

Storage

Storage on the 4D Cloud platform is provided through a dedicated 10 Gbps storage network
giving your virtual machines access to three tiers of storage:

- **SSD:** Extremely fast random read/write IO and ideal for databases. Our tier-1 SSD storage provides the fastest disk performance available. All SSDs are provisioned in RAID-10 arrays with each VM having its own segregated vDisk
- **15k SAS:** Excellent continuous read/write IO with good random seek performance. Our tier-2 15k SAS storage is ideal for the majority of applications which require fast disk access but not extremely random seek times. All 15k SAS disks are provisioned in RAID-10 arrays with each VM having its own segregated vDisk
- **7.2k SAS:** Average continuous read/write IO with average random seek performance. Our tier-3 7.2k SAS storage is ideal for applications that don't require fast disk access or for storing long term content such as images. All 7.2k SAS disks are provisioned in RAID-10 arrays with each VM having its own segregated vDisk
- VMs can be provisioned with multiple vDisks on separate tiers of storage allowing you to separate your data onto the storage that is most suitable for performance and cost

- Each physical hypervisor has 4x 10 Gbps ports to the storage network allowing for a maximum of 40 Gbps throughput to disk ensuring the performance for your VM is there when you need it

Management

Management of your 4D Cloud is provided through our online portal. This gives you complete control of your whole cloud environment. Here are just some of the things you can do:

- Fully control your virtual servers from creation to deletion
- Get full console access to your virtual servers
- Create and manage private networks (vLANs)
- Create firewalls, load balancers and routers with management through the portal
- Take backups, create new templates and restore virtual machines
- Monitor all aspects of your cloud including disk usage, IOPS and server loads
- View detailed logs of all activity within the cloud (Server activity, backups, changes and updates)
- Set limits for your virtual machines or allow them to scale out
- Migrate your virtual machines at will between hypervisors

Auto-Scaling

The 4D Cloud allows you to instantly increase your resources to meet sudden changes in your virtual machines. This means you can start small and easily add CPU, RAM or storage to your cloud as you need it.

You can scale virtual machines up and down (adding/removing RAM, CPU and storage) as well as scale virtual machines out (by duplicating the machine) when limits you set are met.

- **Vertical Auto-Scaling:** Scale virtual machines hosted in your cloud to have more resources automatically when required. This helps maintain performance for the applications you're hosting in your cloud when they are under heavy loads. For example, an application can automatically add more CPU or storage when peaks of user activity occur, and release it again when activity returns to normal
- **Horizontal Auto-Scaling:** Clone a virtual machine in your cloud when it reaches certain limits of utilisation. For example, an e-commerce website might clone itself in
order to balance transactions across more than one server during busy shopping periods. This would help improve availability and performance for online shoppers when your site is busy

Security

Built from the ground up to ensure your servers, data and network traffic are kept private, the 4D Cloud has many layers of security for the platform itself, between users on the cloud, the virtual machines and the network traffic.

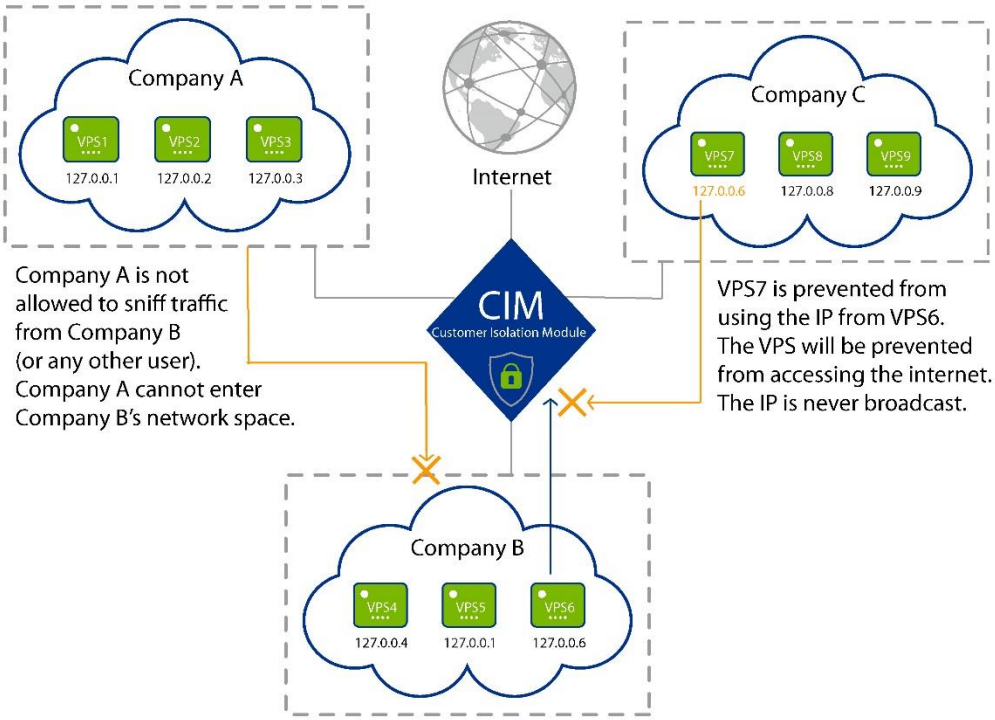

- Four layers of firewalls:
  - Boarder
  - Hypervisor
  - CIM
  - Individual virtual servers
- Private vLANs for network traffic
- Enhanced security templates
- Totally customizable user permissions

You can also deploy your own dedicated firewall appliances within the cloud form our range of pre-configured templates.

Backup

Set up an automated backup for your virtual machines. These can be set to run on a daily, weekly, monthly or annual schedule so you never have to worry about missing a backup again!

Backups can be restored instantly to a virtual machine allowing you to restore your data or applications back to an earlier state. You can also create new server templates from any backup so you can instantly deploy a new server of the same configuration and with the same data.

Open Source and Open Standards

Open Source

The 4D Cloud platform is built upon best of breed open source software utilising the KVM hypervisor with a wide variety of Linux based operating systems available as templates. Software such as Nagios and Cacti is widely used through the organisation. Along with resilient SAN based storage from Infortrend we also utilise the CEPH packages to provide for directly attached, integrated storage within the hypervisors.

Open Standards

Throughout the organisation, 4D uses open standards wherever possible to allow for data interoperability and to prevent ‘lock-in’. The 4D cloud platform has a published API which is available to all users.

Service Management

Compute and Storage Location

4D's head office is based in Byfleet, Surrey along with our flagship data centre (4D Surrey). All services are primary located in this facility with disaster recovery and geographical redundancy provided from our 4D Kent facility.

All support staff are based in the Byfleet office providing 24x7x365 cover.

On Boarding – Inbound Migration

The initial account setup will be done by 4D account managers after a conversation to confirm your compute/storage requirements, any network considerations and managed services that are required.

The initial account will include an administrator user for the cloud portal and the administrator is then able to create sub-accounts for granular user control as well as generate API keys. Predefined user roles are available within the platform including 'billing', 'administrator' and 'support'.

All cloud resources including VM creation, virtual networks (vLANs), firewalls and load balancers can be managed through the online cloud portal or through the API. If specific VMs have been purchased, then these will be pre-configured and available within the cloud portal on your first login. If a pool of resources has been purchased, then resources will be assigned to the user allowing you to create VMs and other appliances up to the limits of your resource pool.

4D has tools available to migrate services and data from existing physical or virtual machines into virtual machines on the 4D cloud platform. These require the installation of a local agent which will image the server before creating a custom template for it within the online portal allowing for an exact replica to be deployed.

For custom deployments the 4D cloud platform allows for ISO images to also be uploaded from which VMs can be booted and installed.

Off Boarding – Outbound Migration

Before any VMs within the 4D cloud platform are deleted clients have the option to either export an image of the VM in the VHD / VMDK formats or store a backup of the VM within the template library for re-deployment at a later date.

4D maintains a secure data destruction process including secure disk wiping and physical disk shredding when appropriate. Any physical media assets are also tracked throughout their life cycle and securely destroyed at the end of client contract.

Backups, Restores and DR

All 4D storage, whether SAN or integrated storage, is based on RAID-10 (data sets are striped along one set of disks then mirrored to another set) which allows for fast I/O performance while remaining fault tolerant.

VMs are able to hot or cold migrate between hypervisors in the event of a physical host failure and all network routes (both external Internet facing and internal storage) are dual path as a minimum across separate physical switches. The 4D storage network runs on multiple 10 Gbps interconnects utilising multi-path IO (MPIO) for iSCSI.

As standard 4D provides for a daily rolling snapshot of your VM which is managed through the online portal. Snapshots can be restored to VMs and manual snapshots can be taken at any time. A custom snapshot schedule can also be configured if longer retention of backups is required (i.e. Grandfather, father and son backups). 4D also provides file based off-site storage on which backups can be stored.

Within the 4D cloud platform clients are responsible for ensuring that backups are completed successfully and 4D is not liable for any failed backups / restores or data corruption.

For high availability applications with specific availability requirements or recovery time objectives (RTO) geographically diverse cloud solutions can be provided with failover between sites or active load balancing between sites.

Boundaries of Service

As part of the public cloud platform 4D is responsible for the hardware health, connectivity and general availability of the platform. Physical hosts are monitored by 4D along with the online portal, core network devices and storage with pro-active

rectification of any faults. As part of the service basic ping checks on VMs can be configured to ensure they are online and responding.

If further support is required then 4D provides a range of advanced monitoring and management services depending on client requirements. These can include pro-active monitoring of web applications, health of individual VMs and operating system / key application monitoring. 4D engineers are also able to provide configuration assistance as well as management of the base operating systems including patch and security management.

The client will be responsible for the deployment, configuration and management of any software/applications/services running on top of the VMs as well as the management of the cloud portal account, users and API.

Levels of service

Infrastructure (IaaS)

4D provides and manages the underlying infrastructure for the cloud platform along with access to the online portal for self-management. 4D will support the virtual machine provisioning process as well as any issues related to the underlying infrastructure but will not provide support for VMs, firewalls, load balancing, network configuration or applications running on them.

Supported Infrastructure (Supported IaaS)

4D provides and manages the underlying infrastructure for the cloud platform along with access to the online portal for self-management. 4D will support the virtual machine provisioning process as well as any issues related to the underlying infrastructure. In addition, 4D engineers are available 24x7x365 to assist with individual VM issues, firewall configurations and network configuration.

Managed Infrastructure (Supported IaaS)

4D provides and manages the underlying infrastructure for the cloud platform along with access to the online portal for self-management. 4D will support the virtual machine provisioning process as well as any issues related to the underlying infrastructure. In addition, 4D engineers will pro-active monitor VMs, web services and applications as well as providing support and management for the base operating systems (patching, configuration and security).

Management Information

Clients are able to access real-time information on their cloud platform through the online portal. Authorised users are able to view the following management information:

1. **Current server status**
   - Powered on / off
   - Disk utilisation
   - CPU utilisation
   - RAM utilization
   - Activity logs
2. **Backup status**
   - Backups which are running
   - Time backups were taken
   - Status of backup (Successful or failed)
3. **Bandwidth**
   - Current bandwidth utilization
   - Historical bandwidth utilization
4. **Overall cloud performance**
5. **Current firewall rules**
6. **Load balancer status and configuration**

If you have taken one of our monitored or managed packages then you will also be able to view detailed information on the following:

- Services and ports (i.e. MySQL, HTTP, POP3, etc)
- Web application status (i.e. If a web page is up, response times and checks to ensure correct data is being provided)
- VM load and health monitoring
- Application status and availability (i.e. MySQL cluster replication, MS SQL health, IIS health, etc)

Service Level Agreements

4D Cloud comprises two related subsystems, the control panel and the cloud platform, each with its own service level agreement.

1. Control Panel

The “control panel” is the web based interface along with the API (Application Programmable Interface) used to create, manage, and delete virtual machines within the cloud platform.

4D guarantees the control panel will be available 99.9% of the time in any given monthly calendar period.

SLA Credit for Control Panel:

Monthly Availability Credit Percentage*
100% - 99.9% 0%
< 99.9% - 99.5% 10%
< 99.5% - 99.0% 20%
< 99.0% 30%

“Monthly Availability” is calculated by availability of the control panel at https://portal.4dcloud.com, on a per customer basis, for a given monthly billing period, as follows (represented as a
percentage): 1 – (Total Failed HTTPS Request)/(Total Valid HTTPS Requests).

* Credits will be calculated as a percentage of all 4D Cloud monthly charges for the calendar month.

2. Cloud Platform

The “cloud platform” includes a virtual machine created via the control panel plus supporting systems and services required for the proper functioning and availability of a virtual machine. Clients manage and have full control over their virtual machine. 4D guarantees proper functioning of the following supporting systems and services:

Network availability

4D guarantees that the data centre network will be available 99.999% of the time in any calendar month, excluding scheduled or emergency maintenance. Network downtime is measured when a virtual machine is not reachable as a result of a failure

in the data centre network. 4D cannot guarantee routing, latency or packet loss once data traffic has left its own network however 4D, so far as is practicable, will configure its routers and switches to ensure outbound data traffic is routed via the available carriers with the best routes to the destination addresses.

4D shall not be liable for network instability or unavailability arising from a denial of service attack directed at or originating from one or more servers located within a data centre through which network connectivity originates or is routed.

SLA Credit for Network Downtime: 1% of the monthly virtual machine charges for every 1% of network downtime below the guaranteed availability, up to a maximum of 30% of the monthly charges in any given month the guarantee is not met.**

Network Packet Loss

4D guarantees less than 0.2% packet loss at any of its outgoing backbone routers during any calendar month.

SLA Credit for Network Packet Loss: 2.5% of the monthly virtual machine charges in any given month the guarantee is not met.**

Data Centre Infrastructure

4D guarantees that data centre HVAC and power will be functioning 100% of the time in any given calendar month, excluding scheduled or emergency maintenance. Data centre infrastructure downtime exists when a virtual machine experiences a failure as a result of power or heat problems.

SLA Credit for Data Centre Infrastructure Downtime: 1 day of the monthly virtual machine charges for every 1 hour of downtime up to a maximum of 30% of the monthly charges in any given month when the guarantee is not met. **

4D Cloud Hosts

4D guarantees the functioning of all 4D Cloud hosts including the hypervisor. If a 4D Cloud host fails, 4D guarantees that restoration or repair will be complete within one hour of problem identification.

SLA Credit for Cloud Server Host Downtime: 1% of the monthly virtual machine charges for every additional hour of downtime, after the first hour, up to 30% of the virtual machine charges.**

Migration

If a cloud server migration is required because of 4D Cloud host degradation, 4D will notify the Client at least 24 hours in advance of beginning the migration, unless 4D determines in its reasonable judgment that it must begin the migration sooner to protect a virtual machine or data. 4D guarantees the migration will be complete within three hours of the time it begins the migration.

SLA Credit for Migration Downtime: 1% of the monthly virtual machine charges for each additional hour of downtime, after the first three hours, up to 30% of the monthly virtual machine charges.**
** SLA Credits will be calculated as a percentage of the monthly charges for those virtual machines adversely affected by the failure for the calendar month during which the failure occurred.

3. Definitions and Limitations

Definitions

For purposes of the Service Level Agreements outlined above:
“4D” means 4D Data Centres Ltd, the supplier of the cloud platform and associated services; “Failed HTTPS request” is defined as: (i) an HTTP 4xx or 5xx server error response to a valid request or (ii) no response to a valid request because the control panel is down. Network errors or downtime outside of the 4D data centre network do not constitute an error; "virtual machine" means a unique virtual machine instance; "monthly virtual machine charges" means the charges for a Client’s virtual machine or virtual machines for the monthly billing period in which the failure occurred and includes hourly virtual machine usage and bandwidth charges; "4D Cloud host" means the physical server which hosts a cloud server; "data centre network" means the portion of the 4D network extending from the network egress point of a virtual machine to the outbound port of the data centre border router; "power" excludes 4D Cloud host power supplies which is covered as part of the 4D Cloud hosts guarantee; "scheduled maintenance" means maintenance (e.g. repairs, modifications, or upgrades) that is announced at least seven days in advance; “emergency maintenance” means critical unforeseen maintenance (e.g. repairs, modifications, or upgrades) needed to ensure security or reliability; “Services Agreement” means the 4D Services Agreement between the Client and 4D which covers the Services to which this Service Level Agreement relates and capitalised words not defined in this Service Level Agreement shall have the meaning ascribed to them in the Services Agreement and it relevant T&Cs.