DATA CENTRES PROOF OF EVIDENCE

PREPARED BY ASHLEY COLLINS

SECTION 78 APPEAL BY GREYSTOKE LAND

WOODLANDS PARK LANDFILL SITE, LAND SOUTH OF SLOUGH ROAD, IVER

NOVEMBER 2024

CONTENTS

1.0 PERSONAL BACKGROUND

1.1 My name is Ashley Collins. I hold a Masters Degree in Town and Country Planning from the University of Manchester.

1.2 I am a chartered Town Planner, and I am employed as a Planning Consultant at the firm Montagu Evans. I hold the position of Partner and I am based in the London Office.

1.3 I have worked in the private sector as a Planning Consultant since 2008 (16 years). Prior to my appointment at Montagu Evans, I was employed as a Director at Jones Lang LaSalle, who are the leading global provider of real estate and investment management services.

1.4 I have a wide range of experience in all aspects of Town Planning, dealing with both Development Control and Planning Policy. Whilst at Jones Lang LaSalle, I worked closely with the EMEA Data Centre Transactions Team as part of the Data Centre Solutions Group. In this role I was part of the team who provided end to end transactional advice on all stages of the data centre real estate life cycle including acquisition, asset management and consultancy.

1.5 I work primarily in the commercial sector with a particular focus on data centre development. I am responsible for a range of stages in the real estate cycle from site finding and acquisition, site promotion and strategic planning advice. I am also responsible for preparing and submitting planning applications for various scales of development, including hyperscale data centres.

1.6 With regards to the appeal site, in my former role at JLL I was the planning consultant advising on the outline planning application validated on 15th November 2021 (LPA ref. PL/21/4429/OA) which was subsequently dismissed by the Secretary of State (APP/N0410/W/22/3307420). Whilst I did not appear as an Expert Witness at the Public Inquiry, I was retained as an advisor and produced the JLL Technical Note on Need for Data Centres that formed an appendix to the Planning Proof of Evidence (CD.I17). I am therefore very familiar with the site and the proposed development, but the application that is the subject of this appeal was prepared and managed by Mr Hutchison of the Pegasus Group.

1.7 The evidence that I have prepared and provide for this appeal (Ref APP/N0410/W/24/3347353) is true and has been prepared and is given in accordance with the guidance of my professional institution, irrespective of by whom I am instructed. I can confirm that the opinions expressed are my true and professional opinions

2.0 INTRODUCTION

2.1 My Proof of Evidence addresses the matter of the need for data centres at a National level (UK), Regional level (London), and Local level – in terms of the Availability Zone that the proposed development seeks to serve (i.e. the Slough Availability Zones which sits within the London Region). It sets out why a data centre development at this location is appropriate given the very distinct set of operational and business necessities. Availability Zones and Availability Regions are networks of data centres, which were developed by industry leaders including by AWS, Microsoft and Google to provide cloud computing.

2.2 Woodlands Landfill Site is situated within the London Availability Region, currently Europe's largest Availability Region, but which has recently grown at a slower rate than competing European cloud computing centres in Frankfurt and Paris.

2.3 The appeal site at Woodlands sits within the London Region's principal Availability Zone, Slough. The Slough Availability Zone is an internationally recognised location for data centres and is the second largest data hub in the world after North Virginia. The area includes data centre cloud operators such as Cyrus One, Amazon Web Services, NTT, Equinix who, together with the range of other providers, offer extensive network and robust connectivity options including high-speed fibre. Microsoft are also developing their new hyperscale data centre facility in Langley.

2.4 The proposed data centre at Woodlands will provide critical capacity to provide further resilience and reliability to cloud services within the Slough Availability Zone. The site benefits from excellent access to the UK's long-haul fibre network which connects London's data centres to a network of subsea cables that land in the UK. These cables play a crucial role in ensuring high-speed, low-latency connectivity between the UK and other global regions.

2.5 I have been instructed to provide a Proof of Evidence setting out an overview of what data centres are, and identifying the need for digital infrastructure in response to the extraordinary growth to the global datasphere. The Proof also addresses the economic and social demand for digital data and digital transformation, and the role of digital infrastructure in operating even the most basic commerce in the 21st Century. Detailed evidence supporting the economic benefits is being provided to the Inquiry by Mr Stephen Nicol, of Nicol Economics, an expert on economic development and economic need and benefits.

2.6 My Proof also addresses the locational requirements of a functional data centre, and identifies why the appeal site is appropriate for this purpose. It explains this in the context of the multiple Availability Zones within the London Region and demonstrates why data centres located in alternative Availability Zones will not meet the specific need identified in the Slough Availability Zone. This is important in the context of the Secretary of State appeal decision for the same site (CD.I12) when the Secretary of State recognised an identified need for data centres within the Slough Availability Zone, but queried why this need could not be met elsewhere.

2.7 Regarding the topic of possible alternative sites that could meet the need for hyperscale datacentres, this is addressed in the Proof of evidence of David Hutchinson of Pegasus Planning. That Proof addresses the suitability of such sites having regard to critical requirements such as environmental / man-made risk, power, fibre, and proximity within the existing Availability Zone in Slough.

2.8 My Proof demonstrates that there is a pressing need for data centre capacity to meet the overwhelming growth in the digital world driven primarily by cloud computing. It further demonstrates why the development site is appropriate to meet the need within the Slough Availability Zone, and why it is that even if alternative sites were to come forward, that would not negate the need for the appeal site to come forward also as this would be in addition to the appeal site (such is the scale of the need). It will also discuss the unique role of the Slough Availability Zone and the local, regional and global role it serves.

3.0 DATA CENTRES AND HOW THEY OPERATE

WHAT IS A DATA CENTRE

3.1 A data centre is a facility that houses IT operations and equipment, including servers, storage systems, and networking devices. It is designed to support large-scale computing and data storage needs. Data centres are critical for running and managing applications, websites, cloud services, and data storage for businesses, governments, and other entities. Data Centres are the power houses of cloud computing and the 'inter-connectors' of the internet.

3.2 The key components of a data centre's architecture are:

1) **Servers**: The physical or virtual machines that process data and run applications.
2) **Storage Systems**: Devices that store data, such as hard drives, SSDs, and storage arrays.
3) **Networking**: Routers, switches, and cables that enable communication between devices inside and outside the data centre.
4) **Power and Cooling**: Backup power supplies, such as generators and uninterruptible power supplies, and cooling systems to maintain the optimal temperatures for data centres to operate efficiently.
5) **Security**: Both physical security (cameras, guards, access controls) and cybersecurity measures to protect data and infrastructure.

3.3 The primary function of a data centre is to store, process, and transmit digital data. These facilities offer the critical infrastructure needed for advanced digital systems, including space, power, cooling, networking, and security, all maintained under strict environmental controls. Data centres come in various forms and capacities, ranging from hosting a portion of a company's files to supporting an entire IT network.

3.4 According to the 2023 Global Networking Trends Report by Cisco, 92% of organisations use two or more public cloud providers to host their workloads, which is an increase from 82% in 2022. Whilst very few businesses need a dedicated data centre of their own, it is difficult to think of any modern UK business that can survive without the need for data centre amenities and functionality. By utilising off-site purpose-built data centres, businesses can take advantage of superior IT infrastructure, enhanced security, and guaranteed availability of their business-critical data. In fact, many businesses find it more economically viable to manage their IT infrastructure in this manner.

3.5 Each business has unique technology and IT needs, but in most cases, a data centre can provide valuable support. By leveraging the advanced infrastructure of a data centre, businesses can access superior services at a lower cost than what they could achieve independently. This makes data centres a highly effective solution for many organisations, including businesses, individuals, governments, and user groups, to benefit from enhanced technology without incurring prohibitive expenses.

3.6 In assessing the Appeal proposal, it is critical to understand that there are different types of data centre classifications, and each perform an important but differing role in serving businesses and society. Their classification depends on a number of factors, whether they are owned by one or many organisations, how they fit into the topology of other data centres, what technologies they use for computing and storage, the service they provide and even their energy efficiency. The different types of data centre can be summarised as follows:

1. **Enterprise data centres**: These are built, owned, and operated by companies and are optimised for their end users. Traditionally in the same building or very close to the user group, they often grew organically linked directly to the growth of the company and differ greatly in design. The primary focus is uptime rather than energy efficiency or sustainability in design. Note: Uptime is defined as the time during which a computer is operational. Downtime is the time when it isn't operational.
2. **Colocation data centres**: In colocation data centres, a company rents space within a data centre owned by others and located off company premises. The colocation data centre hosts the infrastructure: building, cooling, bandwidth, security, etc., while the company provides and manages the components, including servers, storage, and firewalls. Depending on the scale it may be named "retail" or "wholesale". A retail colocation offers third party data centre space to smaller customer deployments. Retail colocation can have hundreds of customers and dozens of network collections. A "wholesale" colocation data centre is where the data centre operator provides the whole data centre to a single third-party business.

3. **Edge Data Centres:** Smaller in scale (1MW – 4MW) and located at the edge of connectivity coverage. Typically used where time dependent activities require lower latency. The subject of latency will be addressed later in this report (see Availability Zones), in short latency is the measure of delay between two points along a network as data moves through it. These are often used in conjunction with larger data centres which will manage the heavier workloads and data backups.
4. **Hyperscale Data Centre**: These data centres are a very specific form of data centre used by the large technology companies. Hyperscale data centres are typically owned and operated by one company. As the name suggests size or scale is the key differentiator. Typically, the largest form of data centre, a company's deployment may start off relatively small but can scale-up rapidly utilising the hyperscale's redundancy and fault tolerance. Provides at least 40MW of IT Load.
5. **Disaster Recovery:** a facility that organisations use as a backup to ensure business continuity in the event that their primary facility experiences a failure, outage, or disaster. It is designed to replicate and store critical data, applications, and systems, allowing businesses to recover and resume operations quickly with minimal data loss. To avoid being vulnerable to the same regional outage or failure, a disaster recovery data centre is located away from the primary facilities.
6. **Artificial Intelligence -** An Artificial Intelligence (AI) data centre is a facility designed to cope with the enormous power, storage and cooling requirements of AI technology. It contains unprecedented computing resources to handle the resource-intensive training and deployment of complex machine learning models and algorithms. Certain AI data centres are less sensitive to low latency requirements.

3.7 The Appeal proposal is for a hyperscale data centre for public cloud computing. This determines the locational requirements for such development and why it therefore must be situated in the Slough Availability Zone.

3.8 Over the past 65 years, computing infrastructure has undergone three major phases of evolution. The first phase involved a transition from mainframes to central processing units, which were located on-premises and managed by in-house IT teams. The second phase introduced widespread virtualisation, and the development of Enterprise and Colocation Data Centres, which allowed for better resource and increased flexibility in workload management across various computing infrastructures. The third and current phase is marked by the shift to cloud computing and the development of Hyperscale Data Centres.

3.9 The rise of cloud computing has led to the emergence of geographically distributed computing, where data and applications are spread across different systems that are interconnected and integrated through network services, functioning as a unified environment. It has meant the term data centre is now used to refer to the department that has responsibility for these systems irrespective of where they are located.

3.10 Organisations today have the option to build and manage their own hybrid cloud data centres or to utilise public cloud services. As a result, applications no longer reside in a single location; they operate across multiple public and private clouds, managed services, and traditional infrastructures. In this multi-cloud era, the data centre has evolved into a vast and complex ecosystem designed to enhance the overall user experience.

WHAT IS THE CLOUD?

3.11 The "cloud" is a virtual storage space where individuals and businesses can store digital assets such as software, applications, files, and various types of data. While often confused with the internet, the cloud is actually just one part of it, specifically the computing technology that enables access to digital resources housed in data centres. The cloud operates through highly interconnected data centres, often hyperscale, allowing seamless sharing of information and applications without the limitations of physical location. Public and private cloud computing technology enables users to access and share digital resources over networks, providing fast, secure access to data and applications across vast distances.

3.12 In the UK, several cloud service providers are particularly popular and widely used. Amazon Web Services, Google Cloud Platform and Microsoft Azure have the largest market share, with providers such as IBM and Oracle also featuring prominently. There are several third-party cloud service providers in the UK that offer specialised services tailored to different business needs. These providers (such as Alibaba, Rackspace and Salesforce) often focus on managed services, cloud migration, cybersecurity, and other IT solutions. Businesses use these platforms under various leasing terms and receive the benefits of scalability, flexibility and security. These providers often serve businesses looking for more personalised services and strategic IT support, beyond what the major cloud providers typically offer. This is due to the fact that via a cloud platform, data is backed up and secured via multiple locally interlinked data centres, located within a defined geographical radius known as Availability Zones (this is discussed further in paragraphs 3.31 to 3.35).

3.13 The last decade has been defined by the growth of cloud computing with Microsoft, Google, Amazon and others servicing enterprise and government with flexible access to remote servers on-demand and in accordance with their business data processing needs. This transition has offered lower software costs and the ability to scale up and down based on business demands with improved resilience. These organisations, which support 70%-80% of all public and private internet activity today, established their first presence in facilities in Slough and Hayes, leveraging the established fibre network and power backbone in these areas.

3.14 According to Cisco's annual Global Cloud Index, the cloud has become so prevalent that, in 2021, cloud traffic represented 95 percent of total data-centre traffic. This was up from 88 percent in 2016 with the primary reason being the rapid adoption of cloud services by both businesses and consumers. The dual factors of huge increases in the demand for data processing coupled with the transition to the cloud, has resulted in the dawn of hyperscale data centre facilities.

THE ROLE OF HYPERSCALE

3.15 A hyperscale data centre is a facility that is owned and/or operated by the company it supports, such as Microsoft, Amazon Web Services (AWS), Google, IBM, Oracle, Facebook, and Apple. A hyperscale data centre is a large-scale facility designed to efficiently support the needs of massive, cloud-based businesses and applications. These data centres provide extensive, scalable digital infrastructure capable of hosting IT applications and managing large volumes of data. They are designed to serve both individuals and businesses efficiently.

3.16 Hyperscale data centres are often defined by their functionality of interconnectedness, this is necessary to provide high availability and fault tolerance, which in turn provides user's access to their data. Given the global reach of cloud computing, they require massive storage and computing capacity which is at scale in its design.

3.17 As the increased demand for data centre storage and processing grows at a high rate, data centre operators have moved to the hyperscale model because it aligns with the increasing demands for high performance, scalability, energy efficiency, and cost optimisation. Hyperscale data centres are also ideal for supporting technologies like artificial intelligence, machine learning and big data analytics due to their vast processing power and storage capabilities. This trend is also reflected in data centre developers seeking deployments over 50MW.

3.18 There is no one single definitive definition of what is a hyperscale, but the following criteria are key features of a hyperscale facility:

1. **Power capacity**: The typical power consumption in a hyperscale data centres has been defined as an average of 40-50 megawatts (MW) and often grow to 100's of MW.
2. **Number of servers**: A typical definition of hyperscale suggest that they would need a minimum of 5,000 server racks, and often house up to 50,000+ server racks.
3. **Size and Scale**: Scale of hyperscale buildings vary, both in terms of configuration and in lifecycle of the actual development. Within a facility, they can scale up or down rapidly to meet fluctuating workloads without compromising performance. A good example of this is retail operators who see a surge in demand due to a particular sales event.
4. **Energy**: Data Centres use energy to power both the IT hardware and the supporting infrastructure (e.g., cooling equipment). The data centre industry is taking sustainability very seriously, which is reflected in the 2021 Climate Neutral Data Centre Pact. The first of its kind "self-regulatory" mandate intended to drive to Carbon Net Zero by 2030.
5. **Support for Big Data**: They are optimised for processing and storing massive amounts of data, making them ideal for big data and AI applications.
6. **Resilience**: High redundancy and resilience refer to a system's ability to continue operating smoothly even if some parts fail. In a data centre, this means having multiple backups and alternative systems in place, so if one component (like a server, power supply, or network connection) stops working, another automatically takes over without causing any disruption. This setup ensures that services remain available and reliable, protecting against unexpected failures or outages, and minimising downtime for users. Network redundancy might involve multiple connections and pathways to guarantee that data can still be transmitted even if one route is interrupted.

3.19 With respect to resilience, this has manifested into the creation of Availability Zones which are often associated with hyperscale data centres. Availability Zones comprise of a number of stand-alone data centres located within specific regions which are each equipped with independent and redundant power, cooling and networking infrastructure all housed in separate facilities. These interconnected data centres provide the platform in which public cloud services originate and operate. Availability Zones are discussed in more detail in paragraphs 3.31 to 3.35.

3.20 The data centre market has moved towards cloud and hyperscale model, which is driven out of data growth and the need to respond to the underlying large scale of demand. As data traffic has accelerated, the growth in hyperscale doubled over the past ten years. There has been a marked transition from a traditional enterprise model where businesses and people had a personal computer tower unit under the desk to off-site server rooms and now to the cloud.

3.21 Developer operators such as Cyrus One, Microsoft, Google and ARK have known requirements and identified deployments of 50MW plus, rising beyond 100 MW in a single hyperscale data centre. This is a marked shift towards hyperscale deployments, increasing from previous demands of circa 20MW which were more prominent prior to 2020. A key turning point was the Covid Pandemic as society and businesses managed to maintain functionality thanks in part to the advantages of data centre.

3.22 The shift to the hyperscale data centre model has significantly improved energy efficiency in the data centre industry, driven by innovations in design, operations, and technology. One of the primary ways hyperscale data centres achieve energy efficiency is through economies of scale. These facilities are designed to house tens of thousands of servers, allowing operators to spread energy consumption more evenly across a larger number of computing units. By optimising infrastructure, such as power distribution, cooling, and server utilisation, hyperscale data centres reduce overall energy consumption per unit of computing power, making them far more efficient than traditional data centres. This allows hyperscale facilities to achieve lower Power Usage Effectiveness (PUE) scores—a metric used to gauge the energy efficiency of a data centre.

3.23 Data centres need electricity to run their equipment and keep the machines cool. Estimations show that data centre energy use has not significantly changed since 2010, even though the number of internet users has doubled, and global internet traffic has increased 15-fold since, according to the International Energy Agency.

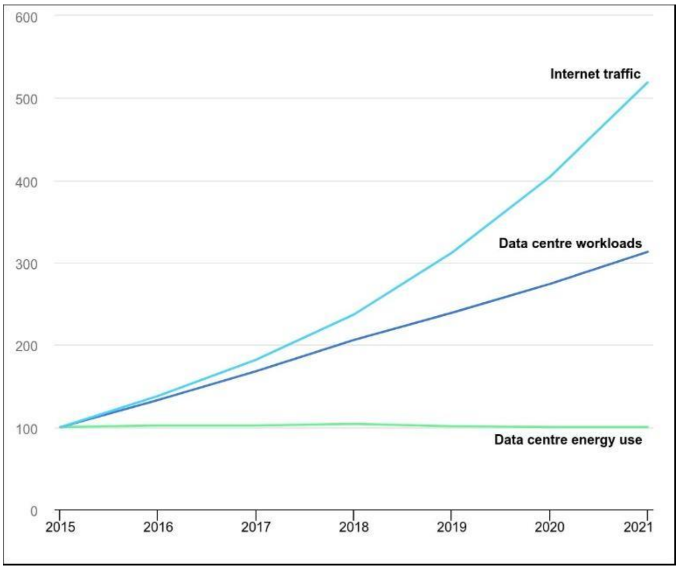

Figure 1: International Energy Agency - Global trends in internet traffic, data centre workloads and data centre energy use, 2015-2021

3.24 Microsoft's white paper 'The Carbon Benefits of Cloud Computing' states that moving IT infrastructures from traditional data centres to the cloud reduced companies' carbon emissions by between 72% to 98%**.** The white paper also concluded that transition to the cloud can result in significant energy efficiency improvements up to 93%. The four main drivers which contribute to the smaller energy and carbon footprint of the cloud-based platform are IT operational efficiency, IT equipment efficiency and data centre infrastructure efficiency which all reduce the energy required to deliver the services. The fourth factor is the purchase of renewable electricity, with many hyperscale operators committed to sustainable energy practices by investing in renewable energy sources. Companies like Google and Amazon have made significant investments in solar, wind, and other renewable energy projects, powering large portions of their hyperscale facilities with clean energy. Microsoft have committed to 100% of their electricity consumption to be matched by zero carbon energy sources by 2030.

LOCATIONAL REQUIREMENTS OF A DATA CENTRE

3.25 While there is a pan-European market (which includes the UK) in the transmission and use of data thanks to the "Cloud", the storage and processing of data in hyperscale data centres can only occur in specific locations. The selection and location of hyperscale and other data centres are shaped by existing infrastructure. Contrary to the common perception, the cloud is made up of physical data centre buildings that are highly connected and resilient. These centres are organised into logical clusters designed to ensure 100% data availability. Their proximity is dictated by the principles of physics, particularly the speed of light, ensuring they are close enough to operate as a seamlessly interconnected network.

3.26 Data centres have very particular locational requirements in relation to existing power and fibre infrastructure and so the location and clustering of existing data centres is no coincidence. They are located where they can be provided with the necessary power (typically with fault independence associated with Grid Supply Points) and other associated fibre infrastructures and where they can meet stringent customer and regulatory requirements.

3.27 The critical location drivers for hyperscale data centres all relate to resilience and business performance. These drivers are:

1. Sufficient size of site to accommodate the scale of facility needed (including expansion space). This scale requirement often leads to a site of the scale that cannot be accommodated within the main urban area of London.

2. Access to an adequate and reliable power supply. The typical power consumption in a hyperscale data centres has been defined as an average of 40-50 megawatts (MW) and often grow to 100's of MW. As the demand for and supply of data centres has grown around London, this is increasingly becoming a constraint on the ability to deliver new data centres in some locations.

3. Access to excellent fibre connectivity. It is important to be close to high-capacity fibre cable networks and to have several diverse fibre routes to ensure resilience.

4. The site needs to be physically resilient - for instance to avoid flooding risk, unstable ground conditions and neighbouring risk exposure such as flight paths, industrial processes and chemical storage.

5. Proximity to other data centres (ideally a minimum of two others) to provide resilience in the event of any failure at the data centre in an "availability zone". Each data centre must be equipped with independent power, cooling and networking.

3.28 Traditional (legacy) data centre architecture design concept was based on a single site location. This provided availability of less than 100% and the overall design relied totally on the resilience of the individual components within a single site. Therefore, if the site or component went down, then availability and access to data was lost as the system did not have in-built resilience or redundancy. The ramifications of this are significant with Gartner reporting that one second of data centre downtime costs institutions **£5,700 per second** - harming both consumers and the economy as well as the firms operating data centres. As an example, the recent CrowdStike outage that affected Microsoft devices is estimated to have cost Fortune 500 companies over $5billion dollars.

3.29 To enhance availability, it became essential to extend reliability beyond individual sites to achieve nearly 100% uptime for users, a critical requirement for cloud services. This need has driven the development of a new architecture in which data centres are both highly resilient and interconnected.

3.30 To enhance resilience and connectivity, the concept of Availability Zones developed. These zones are linked by low-latency networks (measured in milliseconds) to enable near real-time synchronous information transfer and backup. If one data centre or virtual machine fails, another seamlessly takes over the workload without interruption, ensuring 100% uptime and increasing resilience in the cloud architecture.

AVAILABILITY ZONES

3.31 Availability Zones are the key building block of public cloud computing as they allow the Cloud to scale without compromising service quality. Availability Zones are localised clusters of data centres, each equipped with independent power, cooling, and networking infrastructure, located within a defined radius dictated by latency. This proximity is necessary to transfer information instantaneously between data centres and provide near 100% uptime for digital services. Participating data centres in an Availability Zone connect to each other over a resilient, high-speed, low-latency private network link.

3.32 The physical distance between data centres within an Availability Zone is limited by the speed of light to a maximum of 7-10km fibre distance in order to achieve the required sub 2 milliseconds of round-trip latency. This distance is centred at the parent site, which is where a data centre operator first deployed internet capacity in the Availability Zone.

3.33 The distance between sites drives what is what referred to as "latency" or lag. Latency is defined by the Round-Trip Time between sites, which is the time, in milliseconds, it takes for a network request to go from a starting point to a destination, and back again to its starting point. High latency will inevitably lead to network performance issues such as buffering, jitter and packet loss. In contrast, ensuring low latency is essential to enable advanced digital technologies such as 5G wireless networks, artificial intelligence and machine learning (AI/ML) and automation. Availability zones are connected by a high-performance network with a round-trip latency of less than 2 milliseconds.

3.34 Within the Availability Zone, data will continuously and in real-time be transmitted between data centres to carry out processing operations and ensure that services can be provided seamlessly to customers without delay. Moreover, if connectivity at one data centre within the Availability Zone is breached, other data centres within the Zone are able to 'kick-in' and provide continued service delivering near 100% uptime.

3.35 In order to achieve the required sub 2 milliseconds of round-trip latency, this creates a tight radius within which additional sites can be located to allow the expansion of the Availability Zone through the addition of new data centres. If the sites are not within this 2-millisecond round-trip distance, the availability zone is not able to function, services would experience lag and the likelihood of a fault or disruption is high. The primary way to keep latency low is to shorten the distance that data has to travel over the network. Demand for additional services from a particular zone can therefore only be met within the availability zone.

3.36 While all applications suffer some amount of performance degradation due to high latency, applications that rely on human-to-machine interaction will generally be more tolerant of latency than those that rely on machine-to-machine interaction. Many of the advanced digital applications that are in use today must be built at machine speed. Since the potential top speed is so much higher for machines, it's a much greater impact when latency prevents them from reaching it. Modern data centre architecture needs to be able to take in a constant stream of data from diverse sources and perform applications and processes very quickly. The requirement for this process is sub 2 milliseconds and therefore this places a limitation on the distance between the data centre cluster and this is what forms the Availability Zone. It is this limitation of physics which means that a data centre within Manchester serving the Manchester region cannot form part of an Availability Zone in the London Region, such as the Slough Availability Zone, and cannot address the need which resides there. Core Document (CD.K1) details the Microsoft Azure Round Trip latency statistics which they constantly monitor to ensure the network is performing according to requirements and customer needs. The document shows that between UK South and UK West the latency time is 8 milliseconds, thereby demonstrating that sites in the West cannot meet the Availability Zone requirement in the South.

3.37 I note that the draft NPPF has been amended to expressly recognise data centres and to encourage planning policies and decisions to further recognise the specific locational requirements of different sectors and make provision for clusters or networks of data-driven or high technology industries to support their growth (albeit there is already supportive guidance in the existing NPPF as currently drafted). This is discussed by Mr Hutchison in his Proof, but I consider this is an important recognition of Availability Zones in draft national policy. This echoes the recognition contained within the letter from the Department of International Trade (CD.K17) which notes that "*the Thames Valley is central to the UK's data centre landscape, supported by a 21st century digital infrastructure necessary to support data centres*" and that there is "*a* *sustained demand for sites across a corridor that includes Berkshire, Buckinghamshire, Hertfordshire, and west London*". The named locations are reflective of the Availability Zones in Slough, Hayes and Hemel.

AVAILABILITY REGIONS

3.38 The previous section provided an overview of the concept of an Availability Zone, explaining how multiple data centres operate within a single zone by utilising rich fibre and power infrastructure, as well as benefiting from close physical proximity to other data centres. This interconnected setup enables high availability, low latency, and fault tolerance within the zone. This next section summarises the concept of an Availability Region, which consists of a minimum of three Availability Zones working together to enhance resilience and redundancy. Specifically, in the context of the London region, this section will explore how these zones interact within the broader region and the importance of balancing capacity, network connectivity, and redundancy across them to ensure optimal performance and failover capabilities.

3.39 An Availability Region is a geographical location that includes a minimum of three Availability Zones to support cloud services. Every region is isolated (connected to different parts of the electricity grid) and independent from other regions. One region will have multiple Availability Zones, but no Availability Zones are shared among different regions. Regions are spread out all over the world, so cloud providers can reach customers on multiple continents. The London Region therefore has a critical role within the UK, but also as a global service.

3.40 Each major cloud provider operates across multiple Regions around the world, ensuring that services can be provided to customers efficiently, with low latency, and in accordance with local data residency regulations The Proposed Development will serve the 'London' data centre Availability Region. 'London' is the largest of the 'FLAP-D' major markets in Western Europe (which also include Frankfurt, Amsterdam, Paris and Dublin), and is also the largest data centre market in the world outside of Northern Virginia. The table below sets out some of the differences between the Availability Region and the Availability Zone:

Aspect	Availability Region	Availability Zone
Geographical Scope	Large geographic areas (e.g., cities, countries)	Specific locations within a region
Isolation	Isolated from other regions	Isolated from other zones in the same region
Infrastructure	Independent infrastructure	Independent power, cooling, and networking
Redundancy	Multiple zones within a region	Multiple data centres within a zone
Latency	Higher latency between regions	Low latency between zones in the same region
Purpose	Disaster recovery, data residency compliance	High availability, fault tolerance within a region

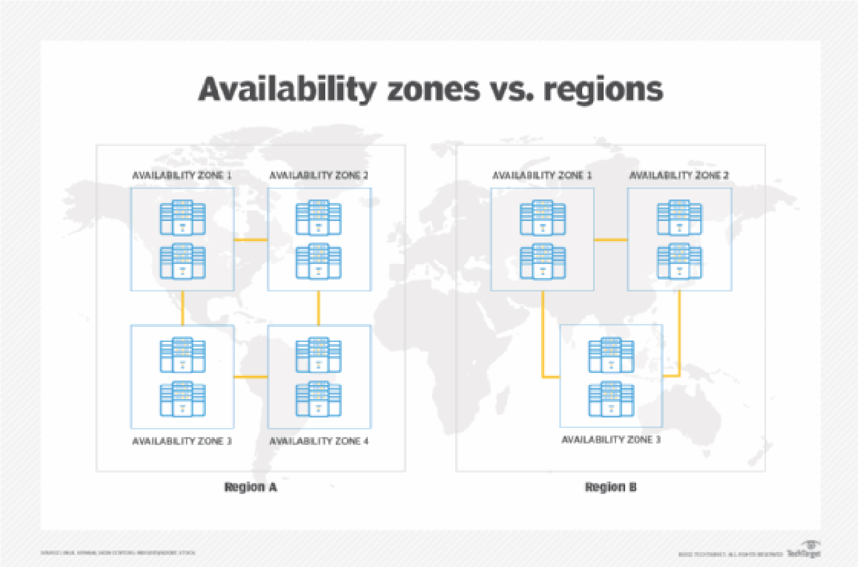

Figure 2: Availability Zones vs Availability Regions. Source AWS

3.41 The practical and geographic implementation of this model is that London is the Availability Region, whilst the separate London Availability Zones are Slough, Hayes, Hemel Hempstead, North Acton and the London Docklands.

3.42 The projected growth of the Slough Availability Zone in combination with the continued growth within the other Availability Zones is necessary to support the wider London Availability Region. Growth within these other Availability Zones does not replace the identified need within Slough, but rather supports this growth and is in addition to this growth.

3.43 Load balancing is a critical element of the cloud computing architecture and is the means by which Availability Zones support each other within a region. Load balancing distributes incoming network traffic across multiple servers or resources. This ensures no single server becomes overwhelmed, which enhances the performance, reliability, and availability of applications and services. Load balancing can apply both within a specific data centre and across multiple Availability Zones, depending on the architecture and requirements of the application or service. Load balancing distributes requests across different Availability Zones, enhancing fault tolerance and availability, if one Availability Zone experiences an outage or degraded performance, the load balancer can route traffic to another Availability Zone, ensuring the application remains available. Therefore, each Availability Zone must have sufficient capacity to handle incoming workload from another Availability Zone in the event of a failure, otherwise it will become overwhelmed by traffic, leading to degraded performance and even failure. For example, AWS's Elastic Load Balancer (CD.K7) can automatically distribute incoming application traffic across multiple Availability Zones within a region as their architecture across three Availability Zones is balanced.

3.44 It is important to stress that 95% of demand is Cloud related and as such the various Availability Zones in the London Region must continue to grow in order to meet the identified need and these are the only specific areas that meet all the unique requirements of the data centre operators. However, whilst this means that – for example – Slough, Hayes or Docklands need to continue to grow in order to meet the need and demand, this is not instead of, or at the expense of another Availability Zones. In this respect, I agree with the evidence of David Hutchison that the need exists in the Slough Availability Zone, and this should be the limits of the search area. Sites outside of this area do not meet the requirements of the Availability Zone architecture as they exceed the sub 2 millisecond latency. Growth within other Availability Zones is necessary, but a data centre within an alternative Availability Zone can only be meeting the need that exists within that Zone.

3.45 It is imperative in the UK national interest that land is made available in the right locations to meet the needs, demands and requirements of the market as this is the location at which the market interest is directed. The creation of Availability Zones has been driven by market demand, leveraging existing infrastructure to support the growing need for resilient public cloud services. As businesses increasingly adopt cloud technologies, the market has responded by developing Availability Zones to ensure high availability, fault tolerance, and low-latency performance. This response is replicated across the globe with London the premier location in Europe. Companies are capitalising on advanced telecommunications, data centre infrastructure, and energy resources to expand these established zones in the London region. The push to grow Availability Zones reflects a market-led response as these are the locations that operators need to locate.

3.46 In this respect I note that the concept of Availability Zones is not a quirk of the UK Planning system, but rather a necessary requirement of cloud computing which is dictated by physics, and replicated globally. For instance, Amazon Web Services currently operate 34 Availability Regions and 108 Availability Zone globally, in Europe this includes Availability Regions in Frankfurt, London, Paris and Dublin, with a new Availability Region in Amsterdam recently announced. Microsoft have 34 Availability Regions and 126 Availability Zones operational; in Europe this includes Frankfurt, London, Amsterdam, Paris and Dublin. I also note that in the United States, where land availability is less of a constrain, Microsoft operate 10 Regions and 24 Availability Zones.

3.47 London is currently the pre-eminent data centre location in Europe and it accounts for the majority of data centre capacity in the UK. The UK is a globally important data centre market (holding 6% of the world market share), home to the largest datacentre market in Europe (holding 42% of market share) and the world's second-largest commercial cluster. However, other European centres, particularly Frankfurt and Paris, are experiencing rapid growth in the supply of data centres and, to some degree, are catching up London. Indeed, within the JLL EMEA Data Centre Report Q4 2023 (CD.K4) it notes that Frankfurt is the fastest growing core market with 134 MW of new supply added in 2023 and a projected 145MW of new supply forecast for 2024. There is, therefore, a national need to deliver data centre development in order for the UK to stay competitive, and this capacity must be accommodated in the London market/region. If further data centre growth is constrained in the UK, investment will be diverted to other European markets which will severely restrict the UK's ability to meet the need for IT services and place a restriction on economic growth. The impact of this is considered in Mr Stephen Nicol's Proof of Evidence.

4.0 THE LONDON AVAILABILITY REGION

4.1 At a European level, there are five well-established central data centre hubs: Frankfurt, London, Amsterdam, Paris, & Dublin (FLAP-D). Each city has specific availability regions that have developed within them, but they are also connected via subsea cables to provide regional interconnectivity. JLL estimates that together, the communication between the availability region across the FLAP-D markets, fuels 80% of the commerce that takes place in Europe

4.2 The London Availability Region is a key hub within the global cloud infrastructure network, playing a crucial role in delivering highly resilient, scalable, and low-latency cloud services to businesses. Strategically located close to one of the world's leading financial and commercial centres, the London Availability Region is designed to support the needs of businesses that require high-performance computing, data storage, and robust disaster recovery solutions.

4.3 London is a leading financial centre, and the development of data centres was driven by banking and financial institutions based in the City of London and Canary Wharf that were seeking to establish 'disaster recovery' data centres in locations outside of London that would act as a back-up to their existing 'in-house' or more centrally located data storage. During the initial build out of modern specification data centres in the 1990's across Europe, London was one of the early hub markets with the Telehouse's being developed in the Docklands. Over time Slough become a focal point for expansion primarily due to the main fibre backbone from the USA East Coast to London which ran along Bath Road. This establishment of rich fibre connections linking subsea cables to Europe and North America meant that over 30 years London has secured a dominant position among European cities.

4.4 The separate London Availability Zones are Slough, Hayes, Hemel Hempstead, North Acton and the London Docklands. The infrastructure within these zones has taken many years and considerable investment to establish, with the clustering of data centres driven by business requirements which, in the case of Docklands and Slough, have taken place over the past 30 years.

4.5 There are three key drivers for the clustering of data centres within the London Availability Region, which are power, fibre infrastructure and proximity to other data centres. Each of these Availability Zones has the existing necessary infrastructure that enables the efficient development of data centres with interconnectivity to other sites.

4.6 Each Availability Zone is a unique entity and serves particular purposes and markets. I therefore set out below the five Availability Zones in the London Region and the primary function they serve, and also a summary of the history of their origin.

SLOUGH

4.7 The growth of Slough as a global data centre hub is linked to its proximity to London and the key infrastructure that supports the operation of data centres, including fibre optic cables and high voltage power.

4.8 Slough is strategically located to the west of London and in proximity to the Great Western Main Line and the Grand Union Canal; both of these house a network of fibre-optic cables for a number of national and international telecommunications providers that carry vital data to and from London through Slough and connect further westwards to the main sub-sea fibre optic cables that run to the United States via connection points in Cornwall and Bristol (and vice versa) as detailed in Figure 3 below.

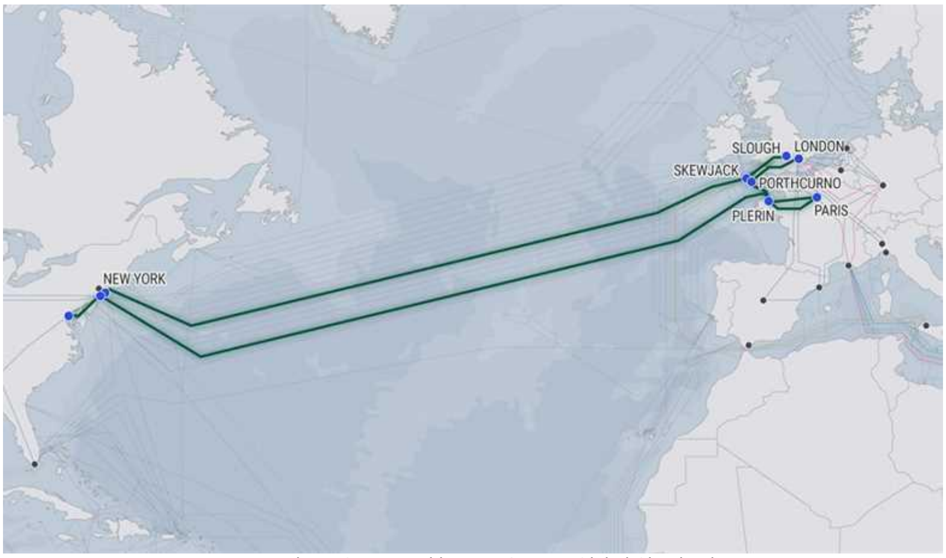

Figure 3: FA-1 cable map, Source: Global Cloud Xchange

4.9 Given the accessibility of Slough to these fibre connections, the area became attractive to banking and financial trading activity, as the area was close enough to London for latency purposes (i.e. minimising any potential delays in the transmission of data related to financial markets or transactions) whilst also giving operators connectivity to overseas institutions in the United States. Slough was therefore able to offer institutions data resilience as part of their emergency and security planning, whilst also providing sufficient land for the development of new, purpose-built data centres at a scale that central London could not.

4.10 Originally founded in 1920, the Estate has expanded over the last 100 years and in more recent years has benefited from the adoption of three Simplified Planning Zones, the first being designated in 1995, renewed in 2004 and the current SPZ last adopted in 2014. The SPZ provides SEGRO, the owners of the Estate, with a number of commercial advantages as it seeks to deliver bespoke premises which serve the needs of modern businesses and provide the sort of facilities that are necessary to continue to attract inward investment and retain its competitive position locally and nationally. The estate also provides much needed digital infrastructure, which together with a progressive and favourable planning regime has enabled growth in tech industries and innovation clusters in Slough.

4.11 The role of Slough as a data centre hub has been further strengthened as a result of the clustering effect of data centres; as banks and financial institutions moved to Slough, infrastructure in the area improved and Slough started to attract supporting technical companies and a work force that could service the data centres.

4.12 This symphony of factors; proximity to major business hub, physical infrastructure, early adoption of technology, favourable planning regime and digital innovation, has resulted in Slough being the second largest data centre hub in the world. Indeed, the data centre hub in North Virginia, experienced similar factors in becoming the largest hub in the world. Ashburn, a city in North Virginia's Loudoun County about 34 miles from Washington D.C, became, not long ago, the first in the world to surpass 1 gigawatt of overall data center capacity.

4.13 The demand for space began to rise as more organisations realised the value of tapping into such a highly interconnected infrastructure leading to the development of more facilities. Today, up to 70% of the world's internet traffic flows through these centres daily which is facilitated by the subsea cable network. Initially all traffic flowed via New York due to the subsea cable (that also linked to London), but in 2017 North Virginia got its first cable linking direct to Europe.

4.14 In summary, the growth of Slough has developed over a 30-year period and has benefited from the development of the required infrastructure over this period of time. The creation of a new Availability Zone to directly match Slough would require significant time and an incredible amount of capital investment and Government intervention. This is not a realistic option.

HEMEL HEMPSTEAD

4.15 Hemel Hempstead's emergence as a data centre hub can be traced back to the early 2000s when the region began attracting interest from tech companies and data centre operators seeking alternatives to the more crowded and expensive areas closer to London, such as Slough. The town's location in Hertfordshire, approximately 24 miles northwest of Central

London, offered the benefits of proximity to the capital while providing more affordable land and a less congested environment for large-scale data centre developments.

4.16 Over the years, companies like Gyron (now part of NTT) and Amazon Web Services have established significant data centre facilities in Hemel Hempstead, leveraging the area's infrastructure and strategic position within the UK's data centre network. The data centres offer robust connectivity and are integrated into NTT's Global network, providing seamless global connectivity. Access to the NTT parent site includes inter data centre connectivity and access to multiple carrier neutral networks, ensuring high availability and low latency for critical operations.

4.17 Hemel Hempstead benefits from a robust and diverse fibre network, providing reliable and high-capacity connectivity to London, other parts of the UK, and beyond. This diverse fibre infrastructure is crucial for the redundancy and low-latency requirements of data centres, enabling them to offer high availability and resilient services.

4.18 The Hemel Availability Zone includes fibre connectivity to the London Docklands and originally provided redundancy requirements to financial institutions. Hemel still caters to demand from some financial customers, but with the presence of NTT and Amazon Web Services has created a significant cloud provision to establish Hemel as an Availability Zone with the ability to link into London and also neighbouring counties. Although it is a smaller Availability Zone within the London Region, the clustering of data centre facilities together with the necessary infrastructure has increased demand for new facilities to provide balance in the workload within the London Availability Region.

HAYES

4.19 The Hayes Availability Zone's popularity as a data centre destination is a result of its strategic location, connectivity to crucial subsea cables, diverse and resilient fibre infrastructure, and robust power availability. Hayes includes large industrial areas which provided easy entry for data centre operators as they fall within the same Use Class and could therefore be repurposed. The transformation into an Availability Zone has largely been driven by the convergence of telecommunications infrastructure, high power availability, and connectivity to international data routes, together with offering a geographic alternative to Slough, whilst benefiting from the same business necessities.

4.20 The presence of diverse and redundant fibre infrastructure in Hayes has also been a significant draw for data centre operators. This diverse fibre network, which includes routes bypassing central London and direct links to other major hubs like Slough, enhances the reliability and resilience of the data centres. Companies like BT, Colt, and other major telecom operators have extensive fibre networks crisscrossing the area, ensuring that data centres in Hayes have access to multiple pathways for both domestic and international data traffic. The availability of dark fibre and the ability to easily connect to other major data centres further adds to its attractiveness.

LONDON DOCKLANDS

4.21 The London Docklands' proximity to the City of London and Canary Wharf, two of the world's leading financial centres, makes it a prime and unique location for data centres. Financial institutions rely heavily on ultra-low-latency data transmission for trading and other operations, and the Docklands offers some of the fastest connectivity available to these critical areas. The London Docklands Availability Zone primarily serves customers in the financial services sector, including banks, trading firms, insurance companies, and fintech businesses.

4.22 The Docklands is also strategically positioned near key landing stations for subsea cables that connect the UK with Europe, North America, and other global markets. These cables, such as the GTT Atlantic and TAT-14, provide low-latency international connectivity, making the Docklands an ideal location for businesses requiring rapid global data exchange.

4.23 One of the key reasons the Docklands has become a major data centre hub is its extensive and redundant fibre optic network. The Docklands is home to Telehouse North, one of the UK's first carrier-neutral data centres, which opened in 1990. This facility, along with others in the area, provides connectivity to a vast array of domestic and international carriers, making it a vital interconnection point for global networks.

4.24 The density and diversity of the fibre infrastructure in the Docklands ensure that data centres have multiple pathways to route traffic, minimising the risk of outages and providing the high availability required by financial services and other mission-critical operations.

4.25 One of the most pressing challenges facing the Docklands Availability Zone is the lack of available space. The Docklands area is highly developed, with limited room for new large-scale or hyperscale data centre facilities. The rapid growth of the financial district has consumed much of the available land. As a result, existing data centres face difficulties in expanding their footprints and increasing density, and new entrants struggle to find suitable sites within the area.

NORTH ACTON

4.26 The North Acton area in London has become an increasingly important hub for data centres. Vantage is currently developing a large campus in North Acton, set to feature two multi-story data centres with a total capacity of 48MW. The campus, which will span 430,000 square feet, is designed to support high-density computing environments with an average power density of 300W per square foot. The first phase of the project is expected to be completed by late 2024. Colt DCS also operates within the North Acton area, providing colocation and connectivity services. Although Equinix has more prominent locations in nearby Slough, it also has operations extending into the North Acton area. The data centres in North Acton primarily serve Hyperscalers, cloud providers, financial services, and gaming due to low latency offering.

SUMMARY

4.27 The London Availability Region is a key hub within the global cloud infrastructure network, playing a crucial role in delivering highly resilient, scalable, and low-latency cloud services to businesses. The separate London Availability Zones are Slough, Hayes, Hemel Hempstead, North Acton and the London Docklands; these together form the London Availability Region. Each Availability Zone is a unique entity and typically serves particular purposes and markets, but all contribute towards the resilience of the London Availability Region.

4.28 Availability Zones are designed to be independent and geographically separate from one another within a specific region. Availability Zones within a region are connected by low-latency, high-throughput networking to provide fast data transfer and replication between zones. The replication between zones is monitored through load balancing to ensure that no single zone becomes overwhelmed and to ensure that the Zones within a Region are able to offer complete fault tolerance. For load balancing to work efficiently, each point of the Availability Region should grow at the similar rate and as there is a growing need for data centres, this growth will be focussed on the established Availability Zones as the most effective and efficient location to direct new growth. If one data centre point lags behind, the overall performance and availability of the system are compromised.

4.29 For these reasons, the identified need that exists in the Slough Availability Zone cannot be served by building a new facility in the London Docklands for example. It also demonstrates that creating additional capacity in the London Docklands, does not mean a reduction of need in another Availability Zone which will continue to need to grow to meet the demand.

5.0 DEMAND FOR DATA AND QUANTIFYING NEED

5.1 The exponential growth in data usage globally is driven by several key forces, including the rise of cloud computing, the proliferation of connected devices through the Internet of Things (IoT), and the increasing digitalisation of industries. As more businesses, governments, and individuals rely on data for decision-making, communication, and operations, the demand for robust infrastructure to store, process, and analyse this data has surged. Additionally, the more recent advent of technologies such as artificial intelligence (AI), big data analytics, and 5G networks has further accelerated data generation and the need for low-latency, high-capacity data centres. This expanding need is reshaping the digital landscape, driving investments in data centres worldwide, and fostering innovation in areas like energy efficiency, connectivity, and security.

5.1 The International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. The Global DataSphere is a service provided by the IDC which quantifies the amount of data created, captured, replicated in any given year. This body of research measures and forecasts the amount of new data created and stored across the globe annually.

5.2 According to the IDC, at the beginning of 2022 it was estimated that the Global DataSphere by 2025 would be circa 175 Zettabytes[1], in December of 2022 the IDC projected that figure to be close to 181 ZB. In 12 months, their projection increased by 6.7 zettabytes. By 2027, the annual volume of data generated worldwide could approach 221 zettabytes and According to the IDC Global DataSphere Forecast, 2024-2028 (May 2024), 394 zettabytes of data will be generated globally in 2028. This shows that the data created, captured and replicated is continuing to grow at higher rates than previously anticipated and therefore the infrastructure required to support that data (data centres) is also increasing. I also note that the IDC forecast has increased since the previous appeal at Woodlands where it was accepted by the Planning Inspector that the need was overwhelming and urgent (paragraph IR120, IR252 CD.I12); that assessment was also accepted by the Secretary of State (paragraph 22 CD.I12). Further, it was suggested by the Council at the inquiry held in respect of that previous appeal that the projected need was just short term and that there was uncertainty beyond the projected period of 2027. The latest projection by the IDC not only shows growth beyond 2027, but a considerable amount of growth with a rate that almost doubles between 2027 and 2028. I acknowledge that this is a global projection, but the London Region has a considerable global market share and it is therefore inconceivable that this global growth projection does not equally apply to London.

5.3 The IDC also predict that worldwide spending on artificial intelligence (AI), including AI-enabled applications, infrastructure, and related IT and business services, will more than double by 2028 when it is expected to reach $632 billion. The IDC advise that the rapid incorporation of AI, and generative AI (GenAI) in particular, into a wide range of products will result in a compound annual growth rate (CAGR) spending increase of 29% over their 2024-2028 forecast period. This compares to the previous estimate in 2022 of 19.2% over the then 5-year forecast period and further strengthens the conclusion that growth in digital infrastructure spending will continue to grow beyond current projections into 2029. This slightly exceeds the Savills forecast which anticipates an annual growth of 25.9% (CAGR) investment spend in 2023 and then an annual growth rate of 15.9% between 2024-2030 (CD.K8). Due to the embryonic nature of AI data centres, I have not accounted for AI need in my forecast of data centre need in the UK. However, I note that the UK Government recognises that the UK is a powerhouse of European AI and is home to a third of Europe's AI start-up companies and twice as many as any other European country (CD.K11). Indeed, the Department for Business and Trade (CD.K3) states that there are over 3,000 AI companies including more than 1,900 dedicated AI businesses in the UK's AI industry. These companies generated £10.6 billion in AI related revenues, employed more than 50,000 people in AI related roles and generated £3.7 billion in Gross Value Added. Whilst I have not accounted for AI data centre need in my growth projections, I am very certain that some of this growth will be directed towards the UK and will be <u>in addition</u> to the need I have set out.

5.4 According to Canalys, global spending on cloud infrastructure services reached US$73.5 billion in Q3 2023, a 16% year-on-year increase. This growth is consistent with previous quarters which shows that the cloud market has entered a stable phase following extreme growth during the Covid-19 pandemic.

[1] A zettabyte is a unit of measurement to describe a computer or other device's storage capacity. One zettabyte is approximately equal to 1 billion terabytes.

5.5 Given this exponential growth in data creation, hyperscale data centres are needed for housing the storage and computing capabilities required to monetise the data and deliver the services we have come to expect and rely upon. Indeed, hyperscale data centres are the most efficient and sustainable way to meet this growth in demand.

5.6 I have demonstrated the necessity of expanded data infrastructure in the significant rise in data generation by consumers and businesses. This increase has been well-documented, illustrating the growing demand for robust and scalable systems to manage and store this data. As we generate more data through various digital activities—ranging from social media use to cloud computing—the need for reliable and resilient data centres becomes increasingly critical.

5.7 In this context, the focus shifts to the UK, particularly the region that hosts the country's key Availability Zones. These zones, which are strategically located in and around London—particularly Slough—play a crucial role in supporting the nation's digital infrastructure. I have explained that these Availability Zones are dictated by the laws of physics and require a sub 2 millisecond round trip in order to achieve the necessary low latency for the services they provide. In this respect I note applications for low latency sensitivity include stock trading, financial transactions, virtual reality, online ad bidding, video streaming, telemedicine and remote surgery, smart manufacturing, smart power grid management, anti-virus software, fraud detection and industrial robotics. These are further summarised in Mr Nicol's Proof of Evidence that sets out key sectors that depend on data services.

THE NEED

5.8 In January 2023, JLL forecast the total demand for data centre capacity within the London Region to increase by from c.2,250MW to 3,100MW over the six-year period 2022 and 2027, with a central forecast of 2,665MW. The forecasts were based upon JLL's then market intelligence on the levels of demand and future need for several of the key global market players over this period to meet the growth in data centre capacity needed which were then extrapolated up to overall market need based on the current market share. JLL forecast this total need to include the data centre markets of cloud, colocation and edge, but it excluded the anticipated need for AI data centres. The full Technical Note is included in CD.I17, but I have set these figures out below. The forecasts for the Slough AZ were based on the then assessed market share of 65% of the London region (which reflected the locational factors and advantages for the Slough AZ I explored earlier).

Forecast annual need for extra data centre capacity in the London market areas (MW of IT load)							
	2022	2023	2024	2025	2026	2027	Total
Upper Range	288	374	459	552	651	758	3,082
Lower Range	210	273	335	403	475	553	2,249
Average	249	324	397	478	563	656	2,666

Figure 4: Forecast annual need for extra data centre capacity in the London market areas (MW of IT load)

Forecast annual need for extra data centre capacity in the Slough AZ (MW of IT load)							
	2022	2023	2024	2025	2026	2027	Total
Upper Range	187	243	298	359	423	493	2003
Lower Range	137	177	218	262	309	359	1462
Average	162	210	258	310	366	426	1733

Figure 5: Forecast annual need for extra data centre capacity in the Slough AZ (MW of IT load)

5.9 The research undertaken by JLL supported the previous planning application at Woodlands in Iver that was subject to a public inquiry and recovered by the Secretary of State (Ref: APP/N0410/W/22/3307420). I note that the Planning Inspectorate and Secretary of State accepted the need figures presented by JLL and concluded that there is no doubt that there is a significant and substantial demand for new data centres (CD:I12). I also note that these figures have been accepted by Buckinghamshire Council within the Statement of Common Ground (CD.H6) as part of this Appeal.

5.10 Following the JLL need forecast that was accepted by Buckinghamshire Council, the Planning Inspector and the Secretary of State in respect of the previous Woodlands Appeal, I also reference the evidence of Stephen Beard of Knight Frank in relation to a planning application at Court Lane in Buckinghamshire that was subject to a Public Inquiry in June 2024 (Ref: APP/N0410/W/24/3337981). Within Mr Beard's Proof of Evidence (CD.K10) he sets out a known requirement to the West of London Availability Zone[2] of 1,700MW within 3-5 years. This assessment of need (1,700MW) is very comparable with the average forecast set out in Figure 5 above (1,733MW).

5.11 I note that the forecast need provided by Mr Beard was also not contested by Buckinghamshire Council at the Appeal.

5.12 I rely upon the previous research undertaken by JLL, but I also refer to research by CBRE (CD.K2), Knight Frank and Savills (CD.K9) who are well respected agents in this sector. Indeed, with regards to Savills (CD.K9), I note that they estimate that European data centre capacity will need to triple by 2027, reaching approximately 22,700 MW to support the growth in demand. Savills assesses total need for extra capacity in EMEA as 22,700MW for the three years 2025, 2026 and 2027, or around 7,500MW each year. Over the same three-year period JLL's forecast is for between 1,431MW to 1,961MW (a central estimate of 1,696MW) of extra capacity needed in London, or just 6% to 8% of this overall assessed total need across the EMEA area as a whole, as forecast by Savills. The UK currently has a global market share of between 5% to 7% according to Synergy Research Group, which shows consistency in the projected figures.

5.13 Therefore, the figures from JLL, together with reference to CBRE, Savills and Knight Frank, form the basis of my assessment of need. The table below utilises the previous JLL figures for the London Region and projects them beyond the previous forecast, this time to 2029, utilising a central growth figure of 20% to reflect historic realised growth in the London Region whilst also recognising that other FLAP-D markets have grown at a faster rate recently (historically between 15-25%).

Forecast annual need for extra data centre capacity in the London market areas (MW of IT load)									
	2022	2023	2024	2025	2026	2027	2028	2029	Total
Upper Range	288	374	459	552	651	758	910	1092	5,084
Lower Range	210	273	335	403	475	553	664	796	3,709
Average	249	324	397	478	563	656	787	944	4,397

Figure 6: Forecast annual need for extra data centre capacity in the London market areas (MW of IT load) 2022-2029

5.14 I have adopted a five-year projection to 2029, as I consider it reasonable to limit forecasts to a manageable period for accuracy. However, I expect growth to continue beyond this timeframe based on current market trends and third party research. As set out in section 5.3 and 5.4, since the JLL figures were produced in 2023, the IDC have revised their Global DataSphere projections and have more than doubled their estimates from 181 zettabytes in 2023 to 394 zettabytes in 2028. The 2023 IDC figure was a factor within JLL's projections, but I have not accounted for their revised increase in the table above. However, I can only conclude that the revised IDC assessment will further increase the identified need. Furthermore, the revised projections by the IDC reject the notion – previously asserted by Buckinghamshire Council at the previous Inquiry – that the identified need is short term and disappears after 2027.

5.15 London saw a recent drop in the amount of new supply (due to supply constraints not a diminution of need) added for the year, with only 73MW added in 2023. To put this into context, this is an 18% drop in the total additional supply added in the year previous and roughly on par with the volumes seen in 2020. This should also be considered in the context of Frankfurt, which saw 134MW of new supply added in 2023 which represented an almost 25% increase.

5.16 London saw the lowest percentage growth in supply out of the core European markets with supply expanding by 8% over the past year. This is a drop in previous years (2018, 2019 and 2021) which saw London achieve a 15-20% increase in supply year on year (CD.K4). I agree with the conclusions of CBRE and JLL that this does not represent a reduction in need or demand, but is instead a result of constraints, such as regulatory restrictions on available and suitable land, or limitations of power supply. The supply and demand imbalance has created an upward pressure on rents with the average colocation rents seeing a yearly increase between 9% - 13%, with London seeing the largest increase in the core European market, further reinforcing that demand remains extremely strong and that supply is a factor impacting rents.

5.17 The new supply of 73MW in 2023 should be considered against the identified need of 273MW (lower) to 374MW (upper) (set out in Figure 6) which therefore represents a shortfall of 200MW to 301MW in 2023 being delivered. This undersupply in the London Region also occurred in 2022 with 90MW of new supply added against an identified need of 210MW (lower) to 288MW (upper) resulting in a shortfall of between 120MW to 198MW. This pattern of undersupply is also set to continue in 2024 as JLL and CBRE are anticipating that between 125MW to 130MW will be brought online (CD.K2 and CD.K4), representing a shortfall of between 210MW to 334MW in the London Region. In the past three years (2022-2024) this pattern of undersupply against the identified need results in a total shortfall of 530MW to 833MW. This previous

[2] It is understood that Mr Beard is referring to the Slough and Hayes Availability Zone when describing the West of London Availability Zone.

undersupply has not been added into my projections but does indicate that the identified future need is very likely to increase further.

5.18 In accordance with good practice, I recognise that there is inevitably a degree of uncertainty as to the precise level of future demand and need, hence I have continued with the range that has been produced, but based upon the research and market intelligence, I also consider it inconceivable that there will not be strong and continued demand beyond a five-year period, especially given the developing backlog of new supply compared to need..

5.19 Based on information on current capacity I referred to above, the Slough Availability Zone accounts for 65% of the total capacity across the whole London Region (albeit some estimations place it higher). Utilising the assumption of a continued market share as at present, this translates to a need in the Slough Availability Zone of around 2,097MW to 2,874MW (with a central estimate of 2,486MW) of additional required capacity between 2024 and 2029. These are set out in Figure 7 below.

Forecast annual need for extra data centre capacity in the Slough AZ (MW of IT load)									
	2022	2023	2024	2025	2026	2027	2028	2029	Total
Upper Range	187	243	298	359	423	493	592	710	3305
Lower Range	137	177	218	262	309	359	432	517	2411
Average	162	210	258	310	366	426	512	614	2858

Figure 7: Forecast annual need for extra data centre capacity in the Slough AZ (MW of IT load)

5.20 As with the total London forecasts there is a degree of uncertainty in such estimates, and I have not included an allowance for the historical shortfall in the past two years (an additional 314MW to 430MW). However, the work points clearly to a very substantial level of demand for new capacity in the Slough Availability Zone area over the period to 2029.

5.21 I note that my projected forecast is grounded in the assumption that the Slough Availability Zone maintains its current market share of 65% of the London Region, and note also that there have been substantial planning permissions in the wider London Region that indicate that other Availability Zones in the London Docklands and North Acton are also growing to support the architecture of cloud computing[3]. However, even if the Slough Availability Zone were to see its market share of new data centre capacity fall to 55%, this would still indicate a very substantial additional need for data centres within the area. I also note that Slough Borough Council have recently adopted a revised Simplified Planning Zone which sets out a range of conditions that some types of development, including data centres, can be built without the need to apply for individual planning permission. This is a readoption of the SPZ which has been in place since 1995 but has been updated to create a more sustainable and high-quality business destination. Slough Borough Council were quoted as stating that the new zone could double the number of data centres over the next ten years. However, I note that the SPZ does not contain a specific target or quota for data centre development. Indeed, the SPZ does detail that one of the key objectives is to ensure that it provides sufficient employment generating uses in order to maintain a diverse and varied economy.

5.22 Addressing this demand in the Slough Availability Zone is essential for the UK to retain its competitive edge and appeal to data centre investments, as well as to uphold our top position in the European digital economy. Due to the magnitude of the growth, hyperscale data centres are the primary way to effectively meet this requirement.

5.23 The projected need for IT capacity represents a need to build multiple hyperscale data centres to support the demonstrated demand. If we assume each hyperscale data centre will have a similar IT Load capacity as the Appeal site (c.100MW), then within the Slough Availability Zone alone, based upon my projection of need, a further 20 to 29 will be needed in the next 5-year period in addition to the appeal site.

5.24 I note that the Statement of Common Ground (CD.H6) section 8.15 sets out that the Appellant's assessment of need at the time of the application was not disputed and it was agreed that there was a need in the period up to 2027, to deliver an increase in data centre capacity of around 1,460MW to 2,000MW (a mid-range or average of 1,730 MW) within the Slough Availability Zone. Even on this basis, a further 15 to 20 hyperscale data centres are required in addition to the appeal site.

[3] As previously stated, this growth is needed to meet the identified need within other Availability Zones and should not be construed as meeting the same need, nor meeting a need instead of the need within Slough. This is additional need to the identified need residing within the Slough Availability Zone.

PERMISSIONS GRANTED

5.25 Montagu Evans have undertaken an assessment of all known data centre applications within the Slough Availability Zone, together with their status. I have excluded previous schemes at Wexham Springs and Link Park, as it is understood there is no prospect of these coming forward for a data centre in the immediate future.

Site	Status	Site Area	Size – sqm	Size – MW	Height (storeys and/or metres)
Court Lane Industrial Estate [**USAA Real Estate**] Court Lane, Iver, Buckinghamshire LPA: Buckinghamshire Council	PL/22/4145/OA (outline) APP/N0410/W/24/3337981 – Appeal decision pending.	5.68 hectares	Up to 65,000 sqm across two buildings	90 – 110 MW	Maximum height of 30 metres
Iver Heath Data Park [**CyrusOne**] Dromenagh Farm, Sevenhills Road, Iver Heath, Buckinghamshire, SL0 0PA LPA: Buckinghamshire Council	PL/24/2130/FA Awaiting decision.	16.59 hectares	69,110 sqm	90 MW	
Sergo Park, Iver [**SEGRO**] Thorney Business Park, Thorney Lane, North Iver, Buckinghamshire LPA: Buckinghamshire / Slough Borough Council	PL/22/1775/FA Permission granted	18.3 hectares	Up to 92,304	Unknown but c.90MW	Maximum height of 25 metres
Ridgeway Distribution Centre [**Oakhills Properties LLC**] The Ridgeway Iver Buckinghamshire LPA: Buckinghamshire Council	PL/23/2521/EIASR (EIA Screening Request)	7 hectares	One data centre at 29,800 sqm	Unknown but c.30MW	Two storeys (main building parapet 19.8 metres; highest exhaust fan 22.5 metres; 24.99 lift over run)
Akzonobel Decorative Paints, Wexham Road, Slough, SL2 5DB [**Equinix Limited**] LPA: Slough Borough Council	P/00072/152 Application for outline planning permission validated 9 October 2024. Awaiting decision.	4.89 hectares	90,614 sqm (GEA)	Unknown At least 90MW	41.5 metres (38.5 metres + 3 metres for back up generator flues)
Akzonobel Decorative Paints facility, Wexham Road,	P/00072/096 (outline) P/00072/108 (reserved matters) (Buildings A and	8.47 hectares	Three data centres at 67,337 sqm (GEA)	100 MW	Two to three storeys Facility A: 23.5 metres Facility B: 23.355

Site	Status	Site Area	Size – sqm	Size – MW	Height (storeys and/or metres)
Slough, SL2 5DB [**Yondr**] LPA: Slough Borough Council	B). Buildings A and B under construction (works commenced in May 2022) P/00072/139 (reserved matters) (Building C). Buildings C approved 14 August 2024.				Facility C: 15 metres (all excluding flues and plant)
200 and 210 Bath Road, SL1 3WE [**SEGRO**] LPA: Slough Borough Council	P/20367/001 Section 106 legal agreement to be concluded.	5.46 hectares	79,830 sqm Full planning permission for 30,130 sqm (GEA) Outline planning permission for 49,700 sqm (GEA)	Up to 90 MW (200 Bath Road 40 MW + 210 Bath Road 50 MW)	Building 1 (200 Bath Road): 18 metres to 25 metres Building 2 (210 Bath Road): 25 metres to 31 metres
Bay 9-13 (building 1) Banbury Avenue, Slough Trading Estate, Slough, SL1 4LH [**Equinix**]	P/20054/001 Planning permission granted on 20 May 2024.	0.98 hectares	15,089 sqm	19.2 MW	Five storeys; up to 40.2 metres

5.26 The table above demonstrates that there is up to 600MW of identified supply being promoted in the short term, even assuming that all outstanding applications are permitted. It is therefore evident that the 600MW identified above is significantly less than the need figure for data centre capacity agreed with Buckinghamshire Council in the SoCG of around 1,460MW to 2,000MW (a mid-range of 1,730 MW) between 2022 to 2027 and less still than the figures of 2,411MW and 3,305MW (a mid-range of 2,858MW between 2022 to 2029).

5.27 In addition to the lack of sites in the pipeline, Mr Hutchison also explains in his evidence that both the adopted and the emerging Local Plans do not identify any sites for Data Centres.

5.28 The appeal site is therefore makes a substantial contribution towards the identified need for additional data centre capacity in the Slough Availability Zone. Further, even if other sites were to come forward, they would not stand as alternatives to the appeal site; they would be needed *in addition* to it.

6.0 THE APPEAL SITE

6.1 I have demonstrated that there is a significant need for a new data centre within the Slough Availability Zone to support the broader London Availability Region. As demand for cloud services and data processing continues to grow significantly, further expansion is needed to ensure the region can maintain the required levels of redundancy, fault tolerance, and scalability and support the digital economy and transition to the cloud.

6.2 I have explained in Section 3 and 4 how the Availability Zones function as part of the London Availability Region in providing interconnected data centres to support businesses. I have summarised the growth in data which is being driven by business and society, through increases in mobile devises driven by IoT (Internet of Things), 5G technology, underpinned by the digital transformation. I have demonstrated in Section 5 the anticipated level of need that exists in the London Availability Region and how this growth needs to be directed towards the Slough Availability Zone in order to accommodate the known growth in data production generally and maintain Slough's position as the premier data centre location in Europe and second only to North Virginia globally.

6.3 Site selection for data centre purposes is multifactorial, and specific conditions need to be met (or avoided) for the data centre to be able to function as intended. There are several precise success factors which if not available make it impossible to operate a data centre. One of these is the adequate supply of stable power. For hyperscale data centres to have the required functionality they also need to be in close proximity to each other, a concept known as clustering. Equally important is to have access to independent services and interconnectedness with dedicated fibre optic links, which in turn are connected to international networks.

LOCATION

6.4 The Appeal site is situated within the Slough Availability Zone which forms a key function as part of the London Availability Region which in turn has a significant role globally. The proposed data centre will configure with the existing data centres within the Slough Availability Zone to form a highly connected cluster and boost resilience. The proximity of these data centres within clusters is driven by physics (speed of light) so that they are close enough to function as an interconnected entity.

6.5 The site subject to the current planning application for a new data centre is located in close proximity to existing facilities with the presence of these established data centres satisfying a key requirement for clustering within an Availability Zone, where multiple data centres operate in tandem to ensure high availability, redundancy, and low-latency performance. By situating the new hyperscale data centre near these parent sites, the proposed development will benefit from existing fibre and power infrastructure, and create a seamless extension of the Availability Zones capabilities. This clustering allows the new facility to integrate into the broader ecosystem to support scalable workloads while maintaining the resilience needed for mission-critical services.

SIZE AND TECHNICAL CONSTRAINTS

6.6 When selecting a site for a hyperscale data centre, the site must be of sufficient size to accommodate the expansive footprint required for a hyperscale facility, including room for server halls, cooling systems, power infrastructure, and future expansion. Additionally, ground conditions must be stable, with suitable soil and foundation quality to support the heavy equipment and structures typical of these developments. Flooding risk is a major consideration, and the site should be located outside flood-prone areas to avoid disruptions and damage. Furthermore, the site must be free from other potential risk factors such as proximity to industrial zones with high explosion or fire risks, ensuring the safety and continuity of operations.

6.7 The site at Woodlands meets these criteria in full and is unburdened by technical constraints which would prevent development.

POWER

6.8 At the time of the previous public inquiry at the Appeal Site (APP/N0410/W/22/3307420), the Appellant had power reserved from the Iver substation. However, this power reservation has now been assigned to another scheme. As such a further application for power has already been submitted to National Grid.

6.9 A global data centre provider has an agreement with Greystoke Land to acquire the site upon the grant of planning permission. This operator has a grid connection secured and will be able to bring this to the site if the appeal is allowed.

6.10 Notwithstanding the end users grid connection agreement, there is a well-established secondary market for power connections and parties will trade their agreement for a grid connection; as is evident by the Appellant's previous grid connection agreement being reassigned to another site.

6.11 It is the case, therefore, that the proposed development can operate upon the completion of construction.

FIBRE

6.12 Slough benefits from a robust and diverse fibre network, providing reliable and high-capacity connectivity to London, other parts of the UK, and to the network of subsea cables. This is the most network dense area in the entirety of the UK and Woodlands offers the ability to an operator to effectively connect to the existing cable infrastructure that is located near the site without the need to dig multiple fibre ducts and broker the required easements.

CONCLUSION

6.13 In conclusion, I consider the proposed site offers an ideal opportunity to develop a cloud based hyperscale data centre campus, contingent upon the grant of planning permission. The fundamental requirements for such a development - high-speed fibre connectivity, protection from flooding, and freedom from other significant risks - have been met in full, making this location ideally suited for a large-scale facility.

6.14 The identified need for additional data centres has been evidenced, and the contribution of this development to both the Slough Availability Zone and the wider Availability Region will be substantial. This need cannot be fulfilled elsewhere and whilst future growth within the other established Availability Zones that form the London Availability Region is needed, this growth is in addition to the need that resides within Slough and not instead of it.

6.15 Given the site's highly desirable attributes and strategic location, the data centre market will readily seize the opportunity to develop it, ensuring the delivery of a cutting-edge hyperscale facility that meets the region's growing demands.

7.0 CONCLUSION

7.1 My Proof of Evidence addresses the matter of the need for data centres at a National level (UK), Regional level (London), and Local Level – in terms of the Availability Zone that the proposed development seeks to serve (i.e. the Slough Availability Zone). It sets out why a data centre development at this location is appropriate given the very distinct set of operational and business necessities.

7.2 A data centre is a facility that houses IT operations and equipment, including servers, storage systems, and networking devices. It is designed to support large-scale computing and data storage needs. Data centres are critical for running and managing applications, websites, cloud services, and data storage for businesses, governments, and other entities. Data Centres are the power houses of cloud computing and the 'inter-connectors' of the internet.

7.3 The "cloud" is a virtual storage space where individuals and businesses can store digital assets such as software, applications, files, and various types of data. While often confused with the internet, the cloud is just one part of it, specifically the computing technology that enables access to digital resources housed in data centres. The cloud operates through highly interconnected data centres, often hyperscale, allowing seamless sharing of information and applications without the limitations of physical location.

7.4 92% of organisations use two or more public cloud providers to host their workloads, which is an increase from 82% in 2022. Whilst very few businesses need a dedicated data centre of their own, it is difficult to think of any modern UK business that can survive without the need for data centre functionality.

7.5 The development at Woodlands is situated within one of London's established data centre submarkets in Slough. The Slough Availability Zone is the largest data centre market in Europe and the second largest globally. The area contains the leading data centre providers such as Cyrus One, Amazon Web Services, NTT and Microsoft, who together offer extensive network and robust connectivity options including high-speed fibre connections.

7.6 Data centres have very particular locational requirements in relation to existing power and fibre infrastructure and the clustering of existing data centres. They are located where they can be provided with the necessary power and other associated fibre infrastructures and where they can meet stringent customer and regulatory requirements.

7.7 The London Availability Region is a key hub within the global cloud infrastructure network, playing a crucial role in delivering highly resilient, scalable, and low-latency cloud services to businesses. The separate London Availability Zones are Slough, Hayes, Hemel Hempstead, North Acton and the London Docklands; these together form the London Availability Region. Each Availability Zone is a unique entity and typically serves particular purposes and markets, but all contribute towards the resilience of the London Availability Region.

7.8 Availability Zones are designed to be independent and geographically separate from one another within a specific region. Availability Zones within a region are connected by low-latency, high-throughput networking to provide fast data transfer and replication between zones. The replication between zones is monitored through load balancing to ensure that no single zone becomes overwhelmed and to ensure that the Zones within a Region are able to offer complete fault tolerance.

7.9 For these reasons, the identified need that exists in the Slough Availability Zone cannot be served by building a new facility in the London Docklands for example. It is also the case that creating additional capacity in Slough, does not mean a reduction of need in another Availability Zone which will continue to need to grow to meet the demand.

7.10 The exponential growth in data usage globally is driven by several key forces, including the rise of cloud computing, the proliferation of connected devices through the Internet of Things (IoT), and the increasing digitalisation of industries. As more businesses, governments, and individuals rely on data for decision-making, communication, and operations, the demand for robust infrastructure to store, process, and analyse this data has surged.

7.11 Given this exponential growth in data creation, hyperscale data centres are needed for housing the storage and computing capabilities required to monetise the data and deliver the services we have come to expect. Indeed, hyperscale data centres are the most efficient and sustainable way to meet this growth in demand.

7.12 Utilising the conservative view of a continued market share as at present, I have evidenced that there is a need in the Slough Availability Zone of around 2,097MW to 2,874MW (with a central estimate of 2,486MW) of additional required capacity between 2024 and 2029. I have not included an allowance for the historical shortfall in the past two years (an additional 314MW to 430MW). However, the work points clearly to a very substantial level of demand for new capacity in the Slough Availability Zone area by 2029.

7.13 In assessing the need for data centres, I rely upon the previous research undertaken by JLL, but I also refer to research by CBRE, Knight Frank and Savills who are well respected agents in this sector – as well as the more global assessment by the IDC. The research undertaken by JLL supported the previous planning application at Woodlands in Iver that was subject to a public inquiry and recovered by the Secretary of State. I note that the Planning Inspectorate and Secretary of State accepted the need figures presented by JLL and concluded that there is no doubt that there is a significant and substantial demand for new data centres. I also note that the previous 2022 to 2027 figures presented by JLL have been agreed within the Statement of Common Ground. The data centre need projections contained within my Proof of Evidence build upon this assessment and project forwards to 2029.

7.14 The projected need for IT capacity represents a need to build multiple hyperscale data centres to support the demonstrated demand. If we assume each hyperscale data centre will have a similar IT Load capacity as the Appeal site (c.100MW), then within the Slough Availability Zone alone, based upon my projection of need, a further twenty will be needed in the next 5-year period in addition to the Appeal site.

7.15 Addressing this demand in the Slough Availability Zone is essential for the UK to retain its competitive edge and appeal to data centre investments, as well as to uphold our leading position in the European digital economy. Due to the magnitude of the growth, hyperscale data centres are the primary way to effectively meet this requirement.

7.16 I consider the proposed site offers an ideal opportunity to develop a cloud based hyperscale data centre campus, contingent upon the grant of planning permission. The fundamental requirements for such a development have been met in full, making this location ideally suited for a large-scale facility.

7.17 The identified need for additional data centres has been evidenced, and the contribution of this development to both the Slough Availability Zone and the wider Availability Region will be substantial. This need cannot be fulfilled elsewhere, as the Slough Availability Zone needs to continue its growth to maintain its critical role in supporting the Availability Region and the wider global service. The UK economy's reliance upon data infrastructure is critical and the proposal at Woodlands ensures capacity for future demand, enhanced resilience, and balanced workload distribution across the region.

7.18 Given the site's highly desirable attributes and strategic location, the data centre market will readily seize the opportunity to develop it, indeed an end-operator has an agreement in place, ensuring the delivery of a cutting-edge hyperscale facility that meets the Region's growing demands.

MONTAGU EVANS
70 ST MARY AXE
LONDON
EC3A 8BE

WWW.MONTAGU-EVANS.CO.UK

London | Edinburgh | Glasgow | Manchester