

/Whitepaper



An Introduction to Managed Hosting

vi@tel // simply different

Introduction to Managed Hosting

As more and more business is conducted online, you might see the Internet as potentially offering a cost effective way to manage customers, improve communications and increase revenue. However, you're worried that your current in-house or shared hosting environment will not scale to handle more traffic, and you don't have the budget, time or resource to invest in more infrastructure.

Many businesses believe they face this dilemma, but in actual fact are already spending more than they need to on a hosting solution that is still failing to give them the functionality, flexibility and scalability they need.

Here is a typical example. We met with a company that has several marketing departments. Each department manages their own website and each of these websites is hosted by separate providers.

This hosting model had some implicit strengths:

- The solution is highly redundant as the chance of all hosting providers experiencing downtime at once is minimal
- As most of the websites are hosted by the web developers themselves, there is a rapid turnaround time for changes to be made to the sites

This model also had significant drawbacks:

- The web development and hosting costs were spread across several marketing departments' budgets - hiding their true hosting costs
- The small hosting companies were unable to provide 24/7 monitoring and support - which is more acceptable for brochure-ware sites, but the customer wanted to run increasingly complex applications that needed 24/7 availability

This customer wanted to improve the **manageability, availability and financial control** of their hosting infrastructure - which are 3 common drivers for considering a managed hosting solution.

Other common drivers include **budget, time and resource constraints**. For example, small companies with no in-house IT resource will host their systems externally as it's a cost and time effective alternative to employing an in-house team. However, even larger businesses with a skilled in-house IT team are turning to managed hosting because they need to concentrate on core business activities, and the range of skills and infrastructure needed to manage a solution in-house makes it cost prohibitive to do so. Equipment breaks, becomes out of date, and with the need to have everything backed up and fully redundant - you need to have two of everything, which is just not an option for most IT budgets!

Also, as more and more business is conducted on the web, people need to have an **increasing number of complex applications** such as email, order management and billing systems available for employees, partners and customers to access online 24/7. As we move away from brochure-ware sites, it's no longer acceptable for a business to experience downtime as it can mean lost business and credibility. This is another key reason to outsource this function to somebody else.

But **it's extortionate to outsource and you lose control?** Not so. If anything, it can be far more cost-effective to outsource as you don't need to continually invest in the latest equipment and IT skills. Plus, your hosting provider should offer you a better rate because they host many customers' equipment in their data centres, giving rise to economies of scale. Most hosting providers offer a fully managed service at a fixed monthly fee which you agree up front, making it easier for budget management. Some will also provide online reporting and controlled access to the hosting environment for your key personnel, so that you retain control.

What about the **security** of your managed hosting solution? Good hosting providers will have highly secure, environmentally controlled and resilient data centres in which to house your equipment and systems. For example, there should be dual power supplies, backup generators and hardware to give resiliency, and controlled temperatures and fire suppression units to help manage the environment. There must be controlled access to the building and your service should be monitored and managed 24/7 by skilled technicians. Some providers will also provide managed firewalls and data backup services as standard.

What **service guarantees** can you expect? Service guarantees vary greatly and are sometimes not worth the paper they're written on. We discuss this in further detail later on, but suffice to say it's essential that your hosting provider gives you guarantees that meet your business requirements.

How do you make sure you're hosting solution is **scalable**? Managed Hosting is by its very nature a scalable service allowing you to add capacity quickly and easily. Make sure your hosting provider asks sufficient questions of you up front before recommending a solution. Every business has different requirements and it's important that your supplier takes the time to understand your needs, so that they can recommend the most cost effective, functional, relevant and scalable service for your business.

Choosing the right Hosting Solution

How do you know that a managed hosting solution is going to be right for your business? This white paper offers a comprehensive guide to:

1. Evaluating the right hosting model for your business
2. The building blocks that make up a managed hosting environment
3. Service Level Agreements
4. A Managed Hosting Checklist

1. Evaluating the right hosting model for your business

Here are some key considerations when evaluating the different hosting options available to you:

- **Availability** – Do you need to provide guaranteed access to employees, clients and suppliers around the globe 24 hours a day?
- **Performance** – How important is the end-user's experience? Do you need to provide rapid and secure access to a highly complex site or do you have a straightforward brochure-ware site?
- **Resilience** - If things go wrong e.g. a web or database server failure, how will this impact your business and what's the backup plan?
- **Scalability** – will your solution be able to handle your current and future requirements?
- **Security** – what level of security do you need, particularly as more and more business is being carried out over public networks?
- **Flexibility** – if you have peaks and troughs in activity which result in you needing flexible capacity, can your solution provide this at a cost-effective price, or will you be charged for capacity you don't always use?


Comparison of Hosting Options

Hosting Feature	In-house	Shared Hosting	Managed Hosting
Availability	Depends on internet supplier and in-house monitoring	Depends on internet supplier and level of monitoring	High as monitored and managed and alternate supply available in case of failure
Performance	Depends on Internet access capacity	Depends on no. customer's sharing the environment	Non-restricted Internet capacity available and dedicated environment
Resilience	Tailored (expensive)	No	Comes as standard
Scalability	Depends on in-house capacity	Minimum	Maximum
Security	Tailored	Minimum	Maximum
Flexibility	Limited	Limited	Unlimited

2. The building blocks that make up a managed hosting environment

A good hosting provider will give you dedicated equipment – usually brand name servers – on which to host your systems. A typical hosting environment can be broken down into modular blocks as illustrated in the diagram below. These building blocks are common to most hosting environments, but can vary according to the company's size, corporate IT standards, etc.

The Hosting Environment



Key

- Apps – Applications (email, CRM, website, databases)
- OS – Operating Systems (Microsoft Windows, Linux, Unix)
- Hw – Hardware (HP, Compaq and Sun servers)
- LAN – Local Area Network
- WAN – Wide Area Network

Ideally, your front-end, middleware and back-end servers should have similar hardware and operating system specifications, but the type of applications you choose to run will ultimately determine the equipment you use. Make sure your hosting provider is certified to support a diverse range of equipment and operating systems, as they will be responsible for managing and fixing them in the event of failure.

When demands for capacity are low, you can host multiple applications, such as websites and databases on a single server. However, you need to ensure that your hosting environment can scale to meet increasing demands for capacity. For example, you might need to move your website and company databases onto separate servers over time, and have you thought about what infrastructure you need to have in place for backup? A good hosting provider should help design a flexible solution that is based on your current and projected capacity and requirements.

Data Centre Facilities

The Data Centre is the actual facility in which your managed hosting environment is located and should have the following:

Physical Security	Security policy that employs authentication biometrics, video surveillance, onsite security guards and lockable cages for your servers
Redundant Power	Redundant UPS with automatic transfer to permanent onsite generator in the event of power failure
Diverse Fiber Entry	Diverse fiber entry into and out of the Data Centre from multiple carriers
Climate Controls	Environment controls reduce the risk of condensation or overheating
Fire Suppression	Localised, pre-action, dry-charged/ gas-based fire suppression with preferably very early smoke detection

Since power and the operating environment (the data centre facilities in general) are shared by all components that make up your hosting environment, the power-feed to the server(s) should be carefully designed. Servers that make up a layer of redundancy must have separate power sources, so a loss of a single source can never result in the loss of an entire layer. The combined availability figure of the power must be at around 99.99%. The hosting provider must have generators, batteries and a completely separated power-grid (A/B power) inside its Data Centre to meet this level. These elements are commonly found at hosting providers.

Internet Connection

The Data Centre's internet connection provided to you must be high speed and redundant in line with the following parameters:

- Availability of > 99.97%.
- Low latency (delay), which is measured by:
 - Round trip delay between the Data Centre and any European peering point -faster than 40 milliseconds
 - Round trip delay between the Data Centre and the US east coast - faster than 120 milliseconds
- Packet loss – less than 1%

The speed and redundancy of the internet connection will depend on whether your hosting provider runs their own network. Many hosting providers do operate their own

international IP backbone (network), and the connections from the Data Centre into this backbone must be fully redundant and preferably follow separate routes (fiber ducts) to ensure resiliency.

If the hosting provider does not operate their own IP backbone, the Data Centre must be up-linked to at least two different IP transit networks. The connections into these networks between the Data Centre and the transit providers must be fully redundant and preferably follow separate routes (fiber ducts).

To ensure that your end-users can smoothly access your web applications from any place on the Internet, the networks up-linking the Data Centre, whether operated by the hosting provider or acquired from a 3rd party, must have rich peerings in place. These peering agreements are the interconnections via which traffic is exchanged between providers. Since the websites and their users are typically hosted at different providers, the networks must have redundant interconnections to handle the traffic as it crosses from provider's network to another.

Some hosting providers apply a percentile model to calculate the amount of data traffic you 'burst' above (exceed) an agreed bandwidth threshold. Viatel, for example, applies the 95th percentile, which is calculated in the following manner:

- During the whole month, a continuous measurement is made of your data traffic
- At the end of the month, the highest 5% of measurements are discarded (giving you 5% of your burstable bandwidth free of charge per calendar month)
- The highest measured value remaining in the total set of measurements is the 95th percentile value, which you pay at a pre-agreed bandwidth price. This way, you only pay for what you use and are not charged on a monthly basis for excess capacity that you might only use during peak periods.

Front end environment

The front-end environment handles incoming http and https (website) requests into your hosting environment from the Internet. Other Internet facing elements, like smtp/pop-3 for e-mail, IRC for chat and DNS, can also be added.

To reach the overall availability target of 99.5%, the front-end environment doesn't require additional redundancy. However, redundancy and subsequent performance improvement can be achieved by load balancing web application requests across multiple servers, so that failure of a single machine won't cause your websites to go down. This redundancy also ensures that system software upgrades can be carried out during normal working hours without impacting availability. There are different ways to make a server redundant, and Viatel can advise which method will best match your requirements.


Middleware

As web applications become more complex, database-based storage and retrieval of data is usually added as an extra tier behind either the web servers or behind a layer of transactional middleware.

Back end environment

To make your Managed Hosting environment secure, it is advisable to separate your database server(s) from the front-end environment by using a security device such as a firewall. Even though the minimum deployment consists of a single server with sufficient disk space, you can make your environment more resilient by installing a cluster of two servers.

High level Front and Back End Configuration of a Hosting Environment



Building a cluster places a requirement on the storage (hard disks) for the cluster. In a cluster, one server is fully operational and a second server is put on standby. If the main server fails, the backup server must have access to the actual disks with the database data, indices and transaction logs in order to take over. This requires the database disks to be external to both hosts, which can be achieved by using SCSI or a fiber channel based storage solution. SCSI is by far the cheaper solution of the two. Today disk space is provided at very reasonable price levels.

The external disk storage itself must also be redundant; otherwise, a disk storage failure can cause a complete database failure, since both the main and backup servers would lose access to the disk storage. There are typically three ways to achieve this level of redundancy:

- Buy managed storage from a hosting provider. Many hosting providers have a fiber channel based storage area network (SAN) inside their Data Centre and you rent storage for a monthly fee. This flexible model allows you to increase volumes of data quickly and easily.
- Use external (SCSI based) storage units with redundant power and electronics, which can be serviced whilst powered. This array connects to both servers.
- Use 2 external storage units with the same configuration and mirror the disks 1:1 between these two units. This allows one of the two units to fail without causing any downtime in the cluster.

Fail-over in a cluster is not completely transparent. When a fail-over occurs, there is a short outage of the database. Users that access the front-end may get an error message, but the software should be configured in such a way that there is minimal disruption to the user when the backup systems takes over.

Despite the fact that everything in the cluster, including the data storage, is redundant, regular backup of all data is needed to safeguard against data loss from an operator error or data corruption caused by errors in the software components.

The disk-size of the database servers must be at least 2.5 times the size of the data that is stored in the databases itself. The 'excess' room is needed for:

- Installation of software, such as Operating Systems, Databases, Management Software, etc
- Dumps of databases onto the file system. Trying to backup the data files when the database is running will lead to corrupt backups, so the easiest thing to do is to make a regular dump of all data to the file system and then backup these stable files.

Database management services, such as monitoring, creation of databases, capacity management, backup and restore, must be provided by your hosting provider.

Wide Area Network (WAN)

As well as adding more functionality to your applications, you may need to connect your applications to your existing IT infrastructure so that your customers and partners can access key business information on your network.

Using Dedicated leased lines provides a very high degree of privacy, but there is limited flexibility to upgrade bandwidth or change network topology (moving end-points).

Since the hosting environment is, by definition, connected to the Internet, the most flexible way to interconnect the hosting platform to your back end environment is by using an MPLS based IP VPN. It's easy to set up since it just needs an Internet connection, and is sufficiently secure to prevent unauthorised access.

Local Area Network (LAN)

The network to interconnect your servers within the hosting environment is commonly Fast Ethernet based, but Gigabyte Ethernet is supported more and more. Your hosting provider should provide this network and fully manage it for you.

Security

The hosting provider will typically perform all operational activities relating to the security of your solution; including implementing and modifying the security policies you've set.

Apart from the hosting provider's staff, your application maintenance staff and partners will need to access your Managed Hosting environment to carry out application updates and for troubleshooting purposes. Your hosting provider will grant controlled access – check that this does not affect your Service Level Agreement with them.

The hosting environment must be secured at all levels, from physical access to the data centre, right through to the policies and procedures that your hosting provider chooses to implement. As security is integral to the managed hosting service, it will normally be included in the monthly fee. Security options included in your service might include 'hardened build' (configuring 'out of the box' servers to make them secure); ongoing patch and software management and virus scanning.

Backup and Restore

Your hosting provider should regularly backup your data and this process should not affect the performance of your service or cause any downtime. In the event that you do lose important data, your provider should be able to restore either full or critical production data.

As a preventive measure, you can ask your hosting provider to perform regular test restores of an entire server from the backup tape to ensure that their backup process is

robust. For example, you might ask them to back up and restore all your data as part of the implementation process, and then follow this up with a second test run three months later. After this, the backup and restore test process can be performed on an annual basis.

As an added precaution, a copy of your data can be stored offsite to keep it secure in the unlikely event that the data centre is damaged or destroyed.

Capacity Planning

Your hosting provider is responsible for monitoring and managing the usage capacity of your service. They CPU load, disk utilisation, network load and application performance need to be constantly measured to help plan usage upgrades well before maximum capacity is reached. This is essential to ensure the smooth running of your service and to eliminate the risk of service outages due to insufficient capacity.

Monitoring

Your hosting provider should actively monitor your entire infrastructure from the network right through to your applications, in line with strict Service Level Agreements (SLAs).

Reporting

You should have access to real-time online reports on the performance of your service. The majority of hosting providers offer an online reporting dashboard, with an extract of the monitoring system, combined with cumulative graphs (over days, weeks & months). You should be able to report on all the key performance indicators for your service.

Disaster Recovery

Your provider should have a disaster recovery process that will result in your service being restored in the unlikely event that the data centre is damaged or destroyed. The disaster recovery plan must cater for alternative space, equipment and the restoration of the data on the servers from a backup tape.

If the platform goes into disaster recovery mode, it is important that at least the basic service is restored. Check with your provider what this basic level of service entails as they might not be able to restore all elements of your service.

3. Understanding your Service Level Agreement (SLA)

Every company has its own interpretation of what constitutes a comprehensive SLA; but as more and more business is conducted online, unavailable or poorly performing web applications can lose companies business and damage their reputation.

Your provider must design and operate a service that meets your specific needs; guarantees maximum availability, and operates in line with service guarantees that can realistically be met. The Service Level Agreement (SLA) is a crucial document that sets out specifically what service levels will be provided and gives a binding guarantee that these levels will be maintained to your satisfaction. If these SLAs are not met, make sure your provider has procedures in place to compensate you accordingly.

Viatel will work with you to tailor a scalable hosting solution that meets your specific objectives; enabling your company to run web or IT applications to specific service security, availability, reliability and scalability needs. Our approach to designing your

hosting solution involves minimising potential points of failure by using brand name equipment, high-end data centre facilities, resilient networking infrastructures and active fail-over technologies.

There are so many aspects of a hosting service that can be monitored and proactively responded to e.g. memory capacity, processor speed, network performance and even specific application services and thresholds, and you can choose what aspects you want monitored and managed.

You should constantly review SLA thresholds and measurements to ensure they stay in line with the changing needs of your business.

Metrics

A typical SLA for a Managed Hosting solution will be tailored to:

- Outline service availability metrics
- Show system reports
- Detail compensation calculations and payments
- Include approval processes for maintenance and updating of the solution components and availability metrics
- Identify and document mission critical and non-mission critical components and equipment
- Identify and document processes and metrics
- Document change requests and fault resolution procedures

Response times

Managed Hosting solutions will have many components and time lines. You need to specify what is mission critical and what isn't because this information provides the basis for response times. A typical SLA should break down response times into:

- Service affecting faults/non-service affecting faults
- Planned service requests/urgent service requests
- Standard time and time limit for your service provider to notify you
- Standard time and time limit for an engineer to respond
- Standard time and time limit for escalation of faults of service requests

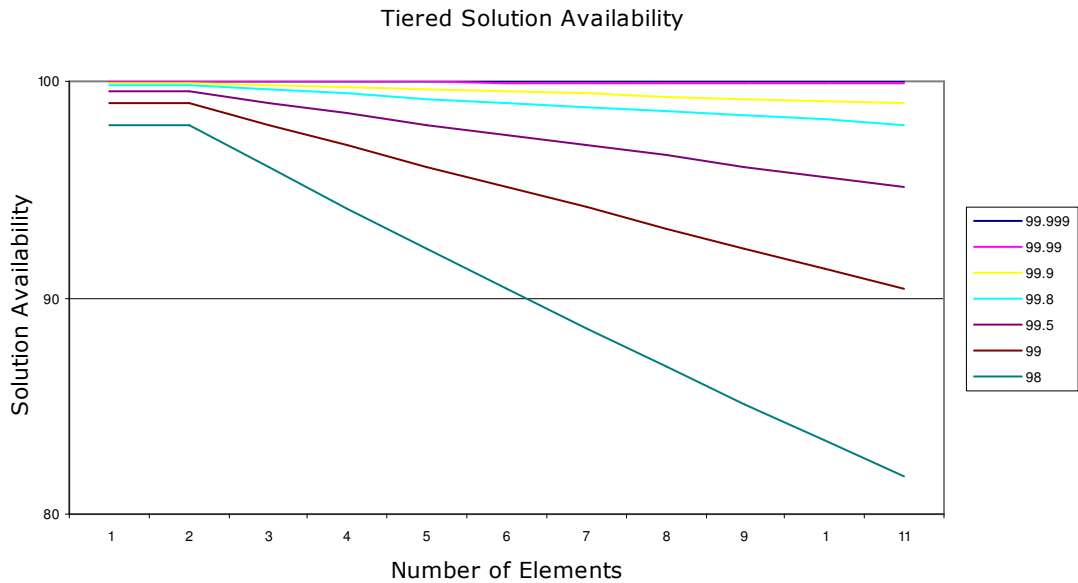
Availability Considerations

A Service Level Agreement (SLA) is about agreeing and managing what is key to your business, and focusing on these elements of the service to ensure they are managed in the appropriate manner.

The SLA content should be tailored to your needs and measurable, with agreed timescales for proactive response (and reactive response if necessary).

Your solution comprises different functional elements (e.g. networks, security devices and servers). You should consider each functional element having its own availability metric; some components have built-in redundancy features, but some simply don't come with built-in redundancy, meaning a lower availability percentage.

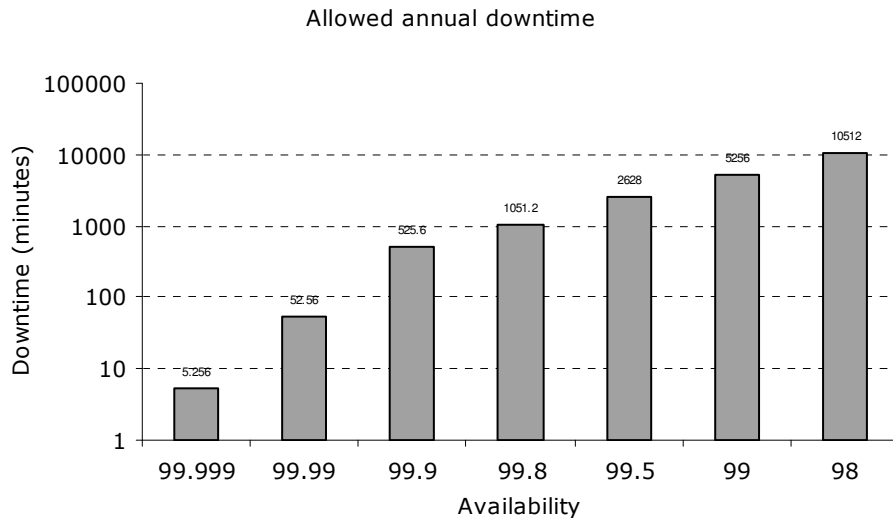
Effect of layering on overall availability



As can be observed from the availability model, the overall availability of your Managed Hosting solution decreases as a function of the number of elements that comprise your solution. Please note that in this model it is assumed that the complete failure of a single functional element will cause a complete loss of service. The effect of failures in the backend can be somewhat limited by building graceful error handling in the web front-end. In a well-designed application, this graceful error handling can be factored in by taking out one functional element in the calculation (e.g. a 5 element application is calculated as a 4 element one).

An overall availability of 98% reads quite high but is actually a very poor figure, since it means having more than one major outage every month (>15 hours outage per annum). As a comparison, the phone and utility networks are managed at the 5-nine (99.999%) level. This constitutes very critical applications. In general, from 99% up, the introduction of every extra 9 will roughly double the costs. Viatel's standard Managed Hosting offering comes with an overall 99.5% availability figure.

Allowed annual downtime as function of system availability



When talking about availability, the mean time to restore (MTTR) and mean time between failures (MTBF) are also very important metrics. In the event of 99.5% availability, the allowed annual downtime is about 43.8 hours (= 3.6 hours per month). There is a huge difference if this manifests itself as a number of short failures or as a single 87.6 hour failure in a 2 year period. It is easier to handle the backlog of short failures than to handle that of one very long failure.

In the full-blown situation, the number of elements (including the network) will be around 4. This means that to get to a total system availability of 99.5%, the 'per element' availability must be at around 99.8% (which is just 54 minutes downtime per month). Since some elements are cheaper to make redundant than others, some fluctuations per element are allowed. In all cases, the 'per element' availability must thus be equal or better than 99.9%. Per functional element availability calculation is required and these must be backed by real live figures from the actual solution once it is in operation.

To gain availability on a functional element that is better than 99.95%, it must have redundancy features built-in. Even the best replacement contracts available in the market are based on 4 hours delivery times, which puts us way out of the MTTR window you want to achieve if hardware replacement is needed to fix a hardware problem. This means that spares must be available on-site or preferably built into the operating platform in the form of clustered servers, double network connections, double firewalls, etc.

4. Managed Hosting Checklist

Here are some key considerations when selecting your hosting provider:

Network coverage

- Does your provider own and manage their own network?
- Does the geographic profile of the network match that of your company?
- Do they have peering agreements with local ISPs in key country markets?

Flexibility

- Do they offer flexible, tailored solutions, as well as off-the-shelf packages?
- Do they have Internet consulting services to assist in the design, build and migration of your solution?
- Do they offer advanced services like application integration and content management?

Reputation

- Do they have a track record of successful implementations?
- Do these implementations include companies of similar size and in a comparable industry to yours?

Support

- Do they have specialised, accredited teams to manage your solution?
- Are these teams available 24 hours a day, 7 days a week? Are they required to respond to customer queries within an agreed amount of time?

Data Centre

- Do they have a network of fully operational data centres across the UK and Europe?
- Are the data centres built with an n+1 configuration so that there are no single points of failure?
- Are they monitored 24x7 by accredited staff?

SLAs

- Do they provide industry leading Service Level Agreements (SLAs)?
- Does the SLA bundle guarantees for all service components into a single contract?

Financials

- Do they have secure financial status?
- Are they continuing to invest in new technologies and resources?

/no problem
getting in touch

For any enquiries about Viatel -
our products and services,
please contact:

Tel: +44 (0) 870 166 2270
Fax: +44 (0) 870 166 2272
Email: sales@viatel.com

Your feedback matters, so if
you have any questions,
comments or suggestions
please email us at:
info@viatel.com

Viatel Holding (Bermuda)
Limited
Inbucon House
Wick Road
Egham
Surrey
TW20 0HR
UK

Tel: +44 (0) 1784 494 200
Fax: +44 (0) 1784 494 201

© Copyright Viatel Holding
(Bermuda) Limited 2004. All
rights reserved. Viatel is a trade
mark of the Viatel group of
companies.

www.viatel.com

Vi@tel // simply different