# DESIGNING A CONCEPTUAL MODEL FOR INTERNET DATA CENTER

SHARMILA DEVI MARIMUTHU

UNIVERSITI UTARA MALAYSIA

2003

# DESIGNING A CONCEPTUAL MODEL FOR INTERNET DATA CENTER

A thesis submitted to the Information Technology School in partial
fulfilment of the requirements for the degree
Master of Science (Information Technology),
Universiti Utara Malaysia

By
Sharmila Devi Marimuthu

**Sekolah Siswazah**
***(Graduate School)***
**Universiti Utara Malaysia**

## PERAKUAN KERJA KERTAS PROJEK
*(Certification of Project Paper)*

Saya, yang bertandatangan, memperakukan bahawa
*(I, the undersigned, certify that)*

**SHARMILA DEVI MARIMUTHU**

calon untuk Ijazah Master of Science (Information Technology)
*(candidate for the degree of)*

telah mengemukakan kertas projek yang bertajuk
*(has presented his/her project paper of the following title)*

**DESIGNING A CONCEPTUAL MODEL FOR INTERNET DATA CENTER**

seperti yang tercatat di muka surat tajuk dan kulit kertas projek
*(as it appears on the title page and front cover of project paper)*

bahawa kertas projek tersebut boleh diterima dari segi bentuk serta kandungan dan meliputi bidang ilmu dengan memuaskan.
*(that the project paper acceptable in form and content and that a satisfactory knowledge of the field is covered by the project paper).*

Nama Penyelia : Dr. Suhaidi Hassan
*(Name of Supervisor)*

Tandatangan :
*(Signature)*

Tarikh : 5/10/03.
*(Date)*

# PERMISSION TO USE

In presenting this thesis in partial fulfillment of the requirement for a postgraduate degree from Universiti Utara Malaysia, I agree that the University Library may make it freely available for inspection. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes maybe granted by my supervisor, in their absence, by the Dean of the Graduate School. It is understood that any copying or publication or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that due recognition shall be given to me and to Universiti Utara Malaysia for any scholarly use which may be made of any material from my thesis.

Requests for permission to copy or to make other use of material in this thesis, in whole or in part, should be addressed to:

**Dean of Information Technology School**
**Universiti Utara Malaysia**
**06010 UUM Sintok**
**Kedah Darul Aman.**

# ABSTRAK

Thesis ini bertujuan membangunkan model konseptual bagi *Intenet Data Center(IDC).* IDC menyediakan kemudahan terhadap kebanyakkan organisasi supaya dapat megukuhkan kemudahan tunggal dan berkonsi kepada *host* pelayan Internet and perkhidmatan pada kadar yang berpatutan. Unsur pembinaan rangkaian IDC adalah berdasarkan beberapa kunci utama seperti pelanggan, *perimeter routers, load balancers, cloned front-end Web servers, multilayer switches, firewalls, infrastructure servers and back-end database* dan pengurusan sistem. Thesis ini juga berfokuskan komponan logikal yang menyediakan kemudahan yang mudah di ukur, memperolehi, selamat dan dapat diuruskan. Kelebihan utama IDC adalah, ia mengurangkan modal dan perbelanjaan pengendalian dalam organisasi. Pemiagaan besar merupakan model kepada perubahan dinamik yang kebiasaannya akan dimulakan dengan permintaan yang kecil and terus membangun. Pembangunan ini dilaksanankan dalam sokongan unik pengguna yang boleh membangun dengan pantas dan juga dalam kerumitan dan integrasi perkhidmatan pelangan yang diberi. Pembangunan ini mesti dibina dalam rekabentuk asas yang kukuh yang dapat menyokong perolahan yang tinggi, inftastuktur yang selamat dan dalam pengurusan infiastruktur.

# ABSTRACT

The purpose of this article is to describe the development of a conceptual model for Internet data center (IDC). IDC is a facility where many organizations can leverage a single, shared infrastructure to economically host Internet servers and services. The key architectural elements of the IDC network include clients, perimeter routers, load balancers, cloned front-end Web servers, multilayer switches, firewalls, infrastructure servers and back-end database and management systems. This paper focuses on the logical components that provide an infrastructure that is scalable, available, secure, and manageable. The major benefits of IDC are, it will reduce capital and operating expenses in an organization. Large businesses are models of dynamic change. They usually start small and grow exponentially with demand. They grow both in the number of unique users supported, which can grow extremely quickly, and in the complexity and integration of user services offered. This growth must be built on a solid architectural foundation that supports high availability, a secure infrastructure, and a management infrastructure.

## TABLE OF CONTENTS

Page

## LIST OF FIGURES

## LIST OF TABLES

# CHAPTER 1

# INTRODUCTION

A data centre is a large data housing infrastructure that provides secure, high bandwidth access to clients for a range of Internet related services. Essentially it comprises servers, firewalls, high bandwidth Internet links and stringent physical security facilities (Dodds, 2000). A data center physically houses various equipment, such as computers, servers (e.g., web servers, application servers, database servers), switches routers, data storage devices, load balancers, wire cages or closets, vaults, racks, and related equipment. Some data centers may have critical requirements for security and reliability anytime.

An Internet data center (IDC) is a subset of data center, which is for all intents and purposes, a warehouse filled with computer servers in a carefully controlled environment. IDCs usually have multiple high-speed Internet connections, just in case one connection goes down, on-site technical support staff, very tight security, and gas-based fire suppression systems (McReynolds, 2001). Dodds (2000) adds that, many large Web hosts have their own data centers, and most of the smaller hosts rent space in them. The data center is (in theory) a safe and secure place for the computer equipment to operate, equipped with security guards, camera, cooling units, generators, guaranteed fuel availability, and support offices.

Large businesses are models of dynamic change. They usually start small and grow exponentially with demand. They grow both in the number of unique users supported, which can grow extremely quick, and in the complexity and integration of user services offered. The business plans for many startups are vetted by their investors for a believable 10-100x scalability projection. Successful businesses manage this growth and change by increasing the number of servers that provide logical services to their clients, either by creating multiple instances of servers (clones) or by partitioning the workload among servers and creating services that integrate with existing computer systems. This growth must be built on a solid architectural foundation that supports high availability, a secure infrastructure, and a management infrastructure (Dodds, 2000).

## 1.1 Problem Statement:

Data Center solutions are complex. Organizations often cannot invest the amount of time and effort it takes to test the configurations and ensure they are fully functional. This results in a solution that is unreliable or unstable. As problems arise short term patches are applied, and at some stage the whole data center becomes an opaque box. The cost of running the system and modifying it to keep up with changing business requirements spirals out of control. Organizations often live with the shortcomings rather than make changes that could break everything. The system cannot adapt fast enough to keep pace with changing business requirements (Dodds, 2000). A typical data center has components that span multiple vendors and multiple products. These components need to be brought together to ensure that they work together as a whole. The components span networking devices, servers, and storage devices. Each of these components has a large number of potentially valid configurations, but only a few of these result in an integrated, functional system. Getting to one of the right configurations can be an expensive undertaking.

In this paper, I will develop a conceptual model to support Internet Data Centre generally.

## 1.2 Research Scope:

In this project, my research scope is limit to:

a) Reviewing the current practice of data centre

b) Study the latest technologies and design a conceptual model of a Internet data center.

## 1.3 Objectives:

a) To study and identify the technical requirements of Internet data center

b) To design a conceptual model of Internet data center.

## 1.4 Research Methodology:

In 1993, the International American Engineering Association (IAEA), offered the following definition for conceptual model

A conceptual model is a set of qualitative assumptions used to describe a system or subsystem for a given purpose. At a minimum, these assumptions concern the geometry and dimensionality of the system, initial and boundary conditions, time dependence and the nature of the relevant physical and chemical processes. The assumptions should be consistent with one another and with existing information within the context of the given purpose.

In this paper, we will be using some methodologies suggested by Microsoft, 2002.

Those are:

a) Identify internet data center service provider and solutions

b) State-of-the-art review (review of the latest technologies)

c) State-of-the-industry review (to review the current practice of the data centre such as MIMOS, Telekom Malaysia, HietechPadu, etc)

d) Determine the evaluation and benchmarking criteria

e) Design conceptual model of internet data center

## 1.5 Significance of Study

Designing a conceptual model of Internet data center will take advantage of innovations in networking, content management, and security to provide IP-based data center applications at new unprecedented performance levels. IDC will reduce capital and operating expenses in an organization. Apart from that, IDC will provide high availability services, instant worldwide presence and superior topological proximity. (Louis, 2002)

## 1.6 Conclusion

This chapter addresses a few essential issues of the study. First and foremost it highlights the current situation and the growing importance of Internet Data Center in today's competitive world. In line with this, this chapter has established the stand for problem statement and the significance of study. Therefore, the objective of this project is to study and identify the technical requirements of Internet data center and to design a conceptual model based on the defined methodologies.

# CHAPTER 2

# LITERATURE REVIEW

## 2.1 Data Center Characteristics

A data center physically houses various equipment, such as computers, servers (e.g., web servers, application servers, database servers), switches routers, data storage devices, load balancers, wire cages or closets, vaults, racks, and related equipment. Some data centers may have critical requirements for security and reliability anytime.

### 2.1.1 Other Names for a "Data Center"

Snevely (2002), quoted that, data center also known as Computer center, datacenter, data storage and hosting facility, server farm, data farm, data warehouse, co-location facility, co-located server hosting facility (CoLo), corporate data center, managed data centers, internet hotel, internet service provider (ISP), application service provider (ASP), full service provider (FSP), wireless application service provider (WASP), telecommunication hotel (or telco hotel), carrier hotel, telecommuncations carriers, High-Density Electronic Loads (HiDEL), or other data networks.

## 2.2 Data Center Design Philosophy

The detailed process of data center design appears on the outset to be a purely mechanical process involving the layout of the area, computations to determine equipment capacities, and innumerable other engineering details. However, the mechanics alone do not a data center make. The use of pure mechanics rarely creates anything that is useful.

There are in fact, some philosophical guidelines that should kept in mind during the data center design process (Snevely, 2002). Some of those philosophies are:

(a) "Look Forward by Looking Back"

(b) "A Modern Pantheon"

(c) "Fundamentals of the Philosophy"

(d) "Top Ten Data Center Design Guidelines"

### 2.3 Benefits of using Data Center

Large businesses are models of dynamic change: they usually start small and grow exponentially with demand. They grow both in the number of unique users supported, which can grow extremely quickly, and in the complexity and integration of user services offered. The business plans for many startups are vetted by their investors for a believable 10-100x scalability projection. Successful businesses manage this growth and change by increasing the number of servers that provide logical services to their clients, either by creating multiple instances of servers (clones) or by partitioning the workload among servers and creating services that integrate with existing computer systems. This growth must be built on a solid architectural foundation that supports high availability, a secure infrastructure, and a management infrastructure. The architectural foundation must meet a number of key design goals.

According to Cassidy (1998), the key benefits for the Internet Data Center include:

- **Scalability**. All components of the architecture must support scaling to provide continuous growth to meet user demand and business requirements.
- **Availability**. Components of the architecture must provide redundancy or functional specialization to contain faults.
- **Security**. The architecture must provide an end-to-end security model that

protects data and the infrastructure from malicious attacks or theft.

- **Manageability**. Ease of configuration, ongoing health monitoring, and failure detection are vital to the goals of high-availability, scalability, and security and must be able to match the growth of the environment.
- **Reliability**. Consistent behavior of system components. Repeatable solutions allow for degrees of predictability when undergoing change and hence reduce risk.
- **Supportability**. An architecture with shared knowledge and intimacy between the customer and all relevant support organizations will lead to supportable solutions with more effective on-going support.

In addition, the solution must provide business value by achieving these goals as efficiently as possible. Wherever possible, without compromising the above design goals, devices used in the Internet Data Center are chosen for cost effectiveness and simplicity. The use of such devices provides the benefit of redundancy without requiring fully redundant equipment. For example, the network switches are configured in such a way as to have all the network traffic balanced across them but they still provide for failover of network traffic.

### 2.3.1 Scalability

Scaling is the ability of a system to handle increasing demands at an acceptable performance level. To achieve scaling, and also to increase security, business Web sites are often split into at least two parts: front-end (client accessible) and back-end systems. Front-end systems generally do not hold long-term state information; instead this is held in back-end data storage. The Internet Data Center scales the number of unique users supported by cloning or replicating front-end systems, coupled with a stateless load-balancing system to spread the load across the available clones. Partitioning the online content across multiple back-end systems allows it to scale as well. A stateful or content-sensitive load-balancing system then routes requests to the correct back-end systems.

The major components that need to be scaled are the network components, front-end Web components, infrastructure/application components, back-end data components, storage components, and management components.

For each component, different dimensions need to scale. For network media it is bandwidth; for Web servers it is processing power; for storage it is size and disk I/O speed.

To scale a system effectively, it is essential to identify the nature of the increasing demand and its impact on the various components. After the component that becomes a bottleneck has been identified, either a *scale-up* or a *scale-out* strategy can be chosen.

Scaling up is a strategy that increases the capacity of a component to handle load. For example, getting a more powerful CPU can scale a Web server. A network component can be scaled up from handling megabytes per second gigabytes per second throughput.

Scaling out is the strategy by which the number of like components is increased, thereby increasing the aggregate capacity of those components. Cloning and partitioning, along with functionally specialized services, enable these systems to have an exceptional degree of scalability by expanding each service independently. For example, the front-end Web can be scaled by adding more servers. Network bandwidth can be scaled by partitioning different types of traffic to different virtual local area networks (VLANs).

### 2.3.2 Availability

Availability is largely dependent on enterprise-level IT discipline, including change controls, rigorous testing, and quick upgrade and fallback mechanisms.

The key to availability is isolating the service functionality from failures of individual components. This can be achieved by removing the dependence, in space and time, of the service on any individual architectural component. Thus, the overall approach for availability is to plan with failures in mind.

Each architectural component of the system is analyzed to verify that it is not dependent on any one piece of hardware performing a specific function or giving access to a specific piece of information. Thus, the architecture requires both redundant components and redundant routing mechanisms, so that requests are always serviceable by a healthy component even in the event of a failure.

Front-end systems are made highly available and scalable through the use of multiple cloned servers which are then load balanced. Back-end systems are more challenging to make highly available, primarily due to the data or state they maintain, but they can be made highly available by using failover clustering for each data partition. Failover clustering assumes that an application can resume on another computer that has been given access to the failed systems disk subsystem. Partition failover occurs when the primary node that supports requests to the partition fails and requests to the partition automatically switch to a secondary node. The secondary node must have access to the same data storage as the failed node, and this data storage should also be replicated. A replica can also increase the availability of a site by being available at a remote geographic location.

### 2.3.3 Security

Managing risk by providing adequate protection for the confidentiality, privacy, and integrity of information is essential to business site success. The key to a successful security implementation is to follow a *defense-in-depth* strategy that defines multiple layers of security and does not rely on any one area to completely secure the infrastructure.

To implement this defense-in-depth strategy the architecture is broken down into separate physical networks or network segments, which allows for the compartmentalization of the system so that a partial compromise of the system does not result in data loss.

The main focus of the security effort lies within two distinct areas:

- Network security
- Host-based security

Network security is generally implemented by breaking up the network into multiple segments and protecting each segment from attack by using various network devices, such as routers with port restrictions, or by using dedicated firewalls.

Host-based security consists of providing each server in the architecture with as much inherent security as possible, so that these hosts do not rely entirely on the network for protection.

A proper security model is crucial within an e-business network, since the perimeter network is exposed to anyone on the Internet. Any e-business site that conducts financial transactions and stores sensitive information, such as credit card data, becomes a target for malicious attacks that can damage a company if the private data is compromised.

### 2.3.4 Manageability

Management and operations broadly refer to the infrastructure, technologies, and processes needed to maintain the health of an Internet application environment and its services. The goals of a management system can be prioritized into the following key areas:

- **Monitoring and alerting**. Keeps track of key events happening in the system and helps identify bottlenecks in the system.
- **Content management**. Allows the system to evolve in a controlled manner as requirements change.
- **Remote management**. Allows the system to be managed from remote locations, which helps to improve system supportability.
- **Backup and restore**. Allows the various systems to be comprehensively backed up, which will then allow any or all systems in the architecture to

restored as required.

There is often considerable synergy between management and the other goals of the Internet Data Center, because an effective management infrastructure provides the tools necessary to meet the other design goals. It is impossible to meet all design goals without an effective management infrastructure, which is why the Internet Data Center relies heavily on these four areas of management.

### 2.3.5 Reliability

Change is the number one factor in a data center that impacts many of the goals that are discussed here. Through consistent behavior of solution components and maintaining a "known working" configuration, the ability to achieve reliable systems that have predictable operation is maximized.

Automated deployment mechanisms are an example of how network, server, and storage components of an architecture can be set up and configured in a repeatable fashion. Such configurations provide the ability to make change in an environment where the starting point and end point are well known, minimizing the risks involved and creating a reliable solution.

### 2.3.6 Supportability

Data center solutions are made up of hardware, software, and service components that come from multiple sources, and none of these sources can know as much about the customer's complete solution as the customer does. Such solutions often put problem isolation emphasis on the customer and make the formulation of customer-centric service level agreements difficult.

The Internet Data Center offers a standardized architecture that, because it is well known to all solution partners, allows them to provide fast and efficient support which can be channeled through a preferred single point of contact. The architecture is a known, supportable baseline; when customized to meet specific customer needs, it can be reviewed for supportability as part of the process of formulating an appropriate support agreement that may be service-level oriented.

## 2.4 Evolution of Data Centers

### 2.4.1 Diversification of Data Center

Internet is expanding not only in the customer base, but also in the spectrum of provided services, which may differ from each other significantly by their requirements of resources and operating modes. The current assortment of data centers, which is presented mostly by Internet data centers serving as Internet service providers, will expand into a broad spectrum of centers that will differ significantly by their scale and by their specialization, starting from application service providers to those centers dedicated to particular E-business segments and even to very specific services. (Kotov, 2000)
For example, the amount of data being stored is growing at an exponential rate, quadrupling every three years. Centralized data storage enables organizations to put essential and frequently accessed data in a single location and make it available to customers anywhere anytime quickly and efficiently. So, storage centers will form the high end of data centers. They will provide high-capacity storage capabilities and specialize at hosting of large "back-end" databases, data warehouses, and large information libraries. The storage centers connected by a network form decentralized storage that nevertheless can provide a single-point view of information for customers through virtualization of these centers.

Mobile data centers provide a provision of back-up systems to the customers in the case of unforeseen circumstances that could interrupt their business operations for an extended period of time. (Rajput, 2000)

Streaming media centers provide the instantaneous delivery of live and on-demand audio, video or other multimedia content over the Internet or enterprise intranets. (Rajput, 2000) They contain presentation servers, archives of past events, data reporting tools, distribution systems that direct audience members to the closest regional center.

Mobile Internet is unique and powerful field for further Internet growth. But it requires seamless integration of multiple wireless networks and mobile devices. Wireless (sometimes also called mobile) data centers act as the wireless providers for companies that don't want to layout the millions it would cost to build wireless centers for themselves. These centers focus on gateway solutions for mobile Internet, converting existing different wireless protocols into others and enabling wireless content providers not to offer their services based on all these protocols.

Internet services, application services, and storage services form a sort of three macro-tiers in the data centers space. (Cassidy, 1998) Different combinations of their capabilities will produce "functionally complete" full service centers that will target different E-service areas and customers (enterprise, electronic delivery services, education, community, etc.).

### 2.4.2 Storage Centers

Information and storage management strategies require special considerations and planning for the following aspects: continuous volume expansion, backup, archiving, data availability, data integrity and security, disaster protection and recoverability, cross-platform data sharing, data manageability.

Gerard (2002) quoted that, the storage capabilities are implemented either as a storage center or as a back-end tier in a data center. In both cases, this is a collection of storage servers that form a storage virtual architecture, on the top of which a storage service infrastructure is implemented.

Modularity is very important from much consideration, primarily due to the volume expansion and the possibility to easy map the services onto the data center architecture.

Storage Area Network (SAN) is a basic paradigm to be developed into the storage centers.

## 2.5 Architecture of Data Center

### 2.5.1 Template Architecture

First of all, we want to outline those features of data centers that are common for all layers and present them as a sort of template architecture of data centers.

Service is a work done by a data center for a hosted virtual data center or for an end customer. Services may be aggregated into hierarchical service suites and consist of subservices. Some other additional relations between the services may accompany the hierarchy relation.

The nature of these relations depends on the specific character of the services. For example, they may reflect actual service dependency: a service may not be completed without requesting the help of another service. Or a service may be called by another service if some special situation is present (an extra service for extra money,). The relations between services force the services to communicate and exchange data.

The frequency of communication and the volume of the exchanged data are important quantitative characterization of the relations. They define the service workload pattern.

The set of all services and all relations between them (together with their quantitative parameters) plus the service workload pattern form the service environment.

Resource is a computing facility (may be a virtual one) that hosts a service or several services and provides the services with everything needed to execute them.

Operating environment maps services onto resources, administer and manage them.

Figure 1 displays a sort of a map of basics notions that we use in the context of the template architecture. Figure 2 instantiate the template architecture for three layers: platform, and e-service layers. Figure 3 shows how different types of data centers can be built using combinations of layers: one virtual center can be built on the top of two platform centers; two virtual centers can be built on the top of one platform center; one e-service center on the right side may combine all layers.

The overall strength and quality of a data center architectural decision will depend mostly on how efficiently its services are mapped onto its resources. (Kan, 2001)

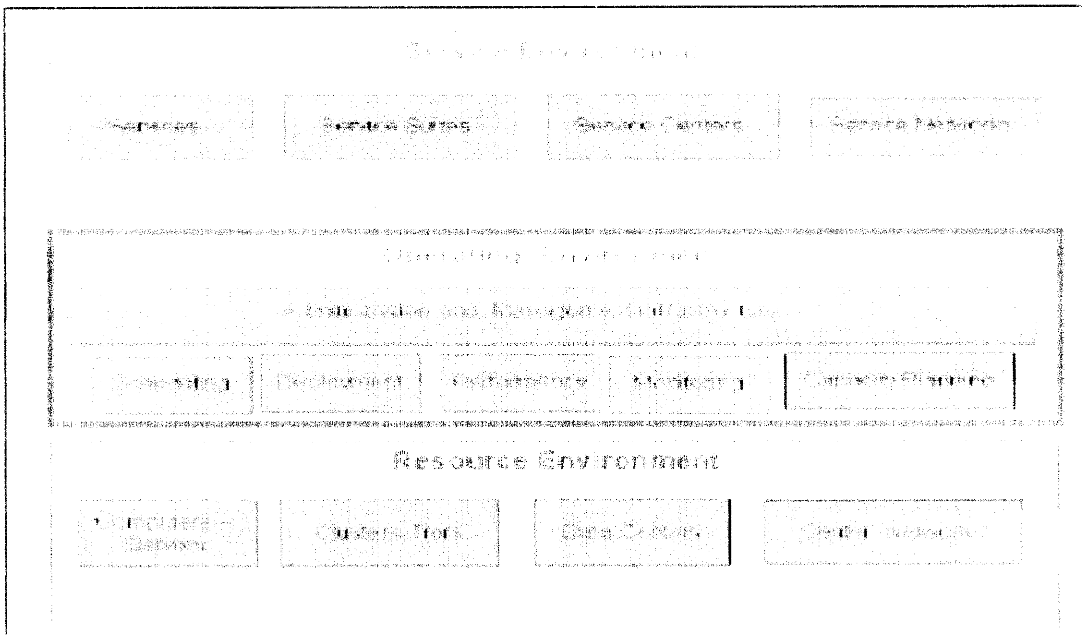

Figure 1.Template data center architecture

*Source: Kan (2001)*

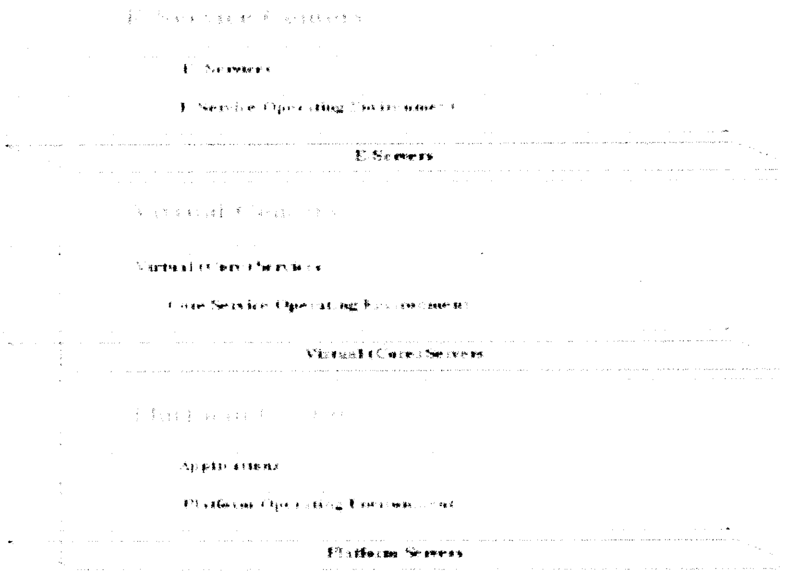
E-Services
E-Service Operating Environment
E-Servers
Virtual Core Services
Core Service Operating Environment
Virtual Core Servers
Applications
Platform Operating Environment
Platform Servers

Figure 2. Three data center layers.

*Source: Kan (2001)*

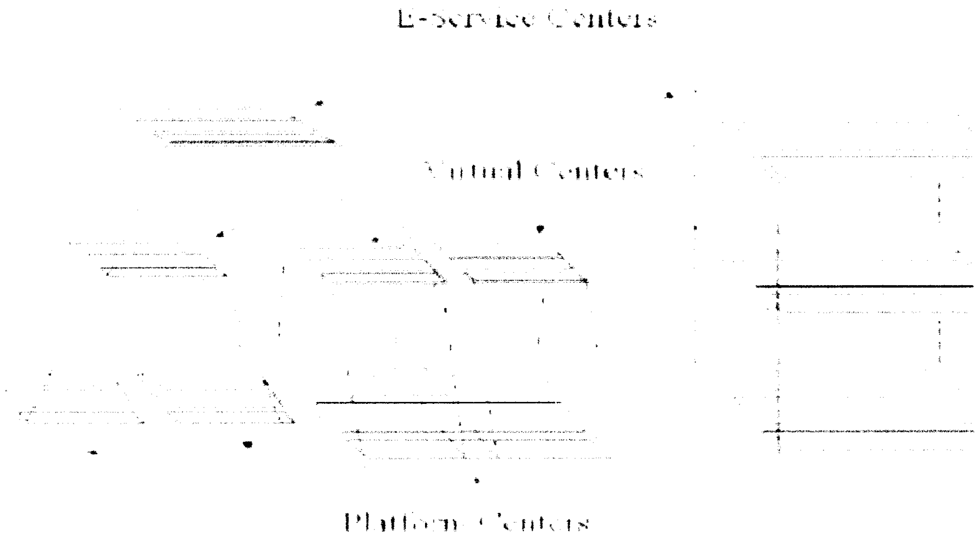

Figure 3. Data centers combined of different layers.

*Source: Kan (2001)*

## 2.6 Operating Environments of Data Center

The life cycle of a data center includes: its deployment; upgrades, modifications, and re-engineering; runtime management. (Goodyear, 1999) An operating environment of a data center should support all the phases of its life cycle and technically is:

- an architecture, including a set of guidelines and standards that define what are the center components (services and resources) and how they fit together and interact,
- a definition for acceptable specifications, interfaces, and protocols both for intra- and inter-center communication,
- administrating of mappings of the center services onto its resources,
- a runtime management environment that facilitate establishing clashless, secure, reliable, and efficient execution of the higher-level tasks and, at the end, the services requested by customers.

Specification of the architectures, interfaces and tasks at all layers should be based on standards (expressed, for example, in XML or some XML-based specifications and protocols).

### 2.6.1 Platform Operating Environment

The main administrative task of the platform operating environment are:
- capacity planning,
- installation and startup,
- resource allocation/reallocation.

Capacity planning is made for virtual data centers hosted by a given platform center, specifically capacity on demand, an incremental planning, when an initial capacity is provided on the customers demand and then is added as needed.

Installation and startup services for software and hardware, as well as software upgrading, versioning, package dependency control, legacy management, are currently costly and time-consuming procedures. Quick deployment of ready-to-use systems and software management (distribution, assignment) will define the quality of services provided by platform data centers.

A platform center that put its resources at disposal of several virtual centers should be able to carry out rapid reconfiguration if some change in its customer population or in customers' demand occurs. Resource allocation includes resource planning, allocation and optimization of CPU time, storage, I/O, bandwidth.

Kan (2001) say that, at the platform level, there are individual or specific management tasks, solutions for which are either available currently or will be available soon, such as:

- network management (connectivity, traffic monitoring and control, diagnostics),
- backup, failure, fail-over and recovery management,
- system evolvement (scaling, rapid capacity re-planning),
- environment and infrastructure management (power supply, external networks),
- monitoring of a large number of system events and dynamic component parameters.

### 2.6.2 Virtual Operating Environment

An important factor of a virtual center performance is a smart organization and control of message and data traffic between the center and its customers and among the center services and servers. (Kan, 2001) An important factor of the traffic efficiency is partition, placement, and replication of services among servers, as it creates a "primary" traffic of requests and responses. The latter generates a "secondary" traffic between services that generate a "ternary" traffic between applications and system programs, and so on. A bad organization of the primary traffic is the main cause of the bad secondary and subsequent traffics.

Main administration problems are:

- Analysis of the e-service parameters (including the e-service workload) in the businesses that will be owners or customers of a virtual data center.
- Analysis of the virtual servers parameters (their capabilities to host services).
- Efficient and secure partitioning (including replication) of services among servers providing an efficient average traffic for statistically dominating workload.
- Automatic (centralized/decentralized) deployment of services on the basis of the partitioning.
- Optimal repartitioning of services onto servers based on the data center statistics.

The dynamic management tasks that are addressed now:

- monitoring of large number of system events and dynamic parameters,
- the non-intrusion measurement of traffic to/from services and servers,
- quality management (availability, guaranties),
- accounting management,
- security and access management.

New management task are related to the service management and to dynamic management of the system load and traffic:

- dynamic repartitioning of data, applications, and services in case of servers failure, uneven load, bottlenecks,
- caching, traffic routing, and load balancing.

### 2.6.3 Operating Environment Structure

It is important to have a proper distribution of the management task between the platform and virtual operating environments. This may be a special problem to resolve, as there is no clear line of demarcation and it may migrate for different types of data centers and in time when some functions of the virtual level will move to the platform level. Being clearly separated, both environments may communicate with each other in a secure way.

Though different in their tasks and administered objects, both operating environments may have similar infrastructure, as well as unified administration and management principles involved.

As data centers are large and complex system, it is extremely important to streamline and consolidate both administration and management tasks in the operating environments:

- reduce the number of data monitored and displayed for the system administrators and managers,
- make all the management tasks feel and look alike, have a common console for the management
- replace gradually the operator-based management by a reactive automatic management of events, load, and traffic,
- decentralize the operating environment; find a right proportion of centralized/decentralized management of the distributed service.

Functions of a dynamic operating environment fall into three categories:

- monitoring and measurement,
- decision making, and
- controlling.

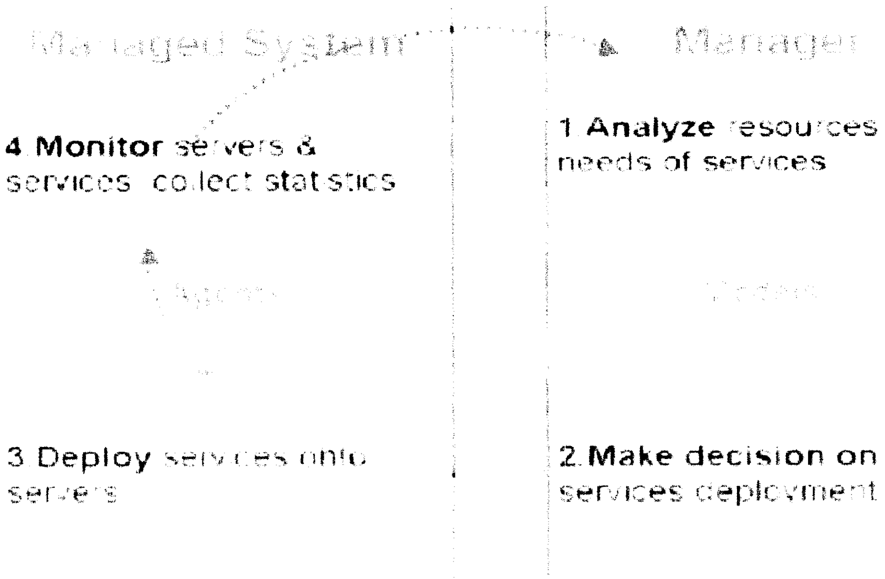
Managed System
Manager
1 Analyze resources needs of services
2. Make decision on services deployment
3 Deploy services onto servers
4 Monitor servers & services collect statistics
Agents
Orders

Figure 4. Management loop.

*Source: Kan (2001)*

Managers and agents implement these functions. Agents do monitoring and measurement and send the monitoring results to managers. The managers make decisions and send orders to agents to make changes in the center structure or/and behavior. To make intelligent decisions, the managers may use data center models that contain information about the center that is necessary and sufficient for finding optimal management solutions. The management loop formed by managers and agents is shown in Figure 4.

Decentralized operating environment of large data centers requires the distribution of management control and responsibility for certain domains to responsible managers.

There may be a hierarchy of managers with managers of lower levels serving also as agents for the higher-level managers.

# CHAPTER 3

# METHODOLOGY

This chapter include discussion regarding the methodology used for this project. The methodologies used for this project are state-of-art review, state-of-industry review and the evaluation and benchmarking criteria.

## 3.1 State-of-art review and state-of-industry review

A compilation of published patent documents to provide either a broad review of a technical field or a detailed report on the most recent developments in a particular area of technology. Analyzed in a time-segmented fashion, this information can be used to display graphically the progress of technological development.

In 1993, the International American Engineering Association (IAEA), offered the following definition for conceptual model
*A conceptual model is a set of qualitative assumptions used to describe a system or subsystem for a given purpose. At a minimum, these assumptions concern the geometry and dimensionality of the system, initial and boundary conditions, time dependence and the nature of the relevant physical and chemical processes. The assumptions should be consistent with one another and with existing information within the context of the given purpose.*

The key architectural elements of the Internet Data Center network include clients, perimeter routers, load balancers, cloned front-end Web servers, multilayer switches, firewalls, infrastructure servers and back-end database and management systems. (Microsoft, 2002) This part focuses on the logical components that provide an infrastructure that is scalable, available, secure, and manageable.

### 3.1.1 Internet Clients

In the Internet Data Center environment, clients issue requests to a service name, which represents the application being delivered to the client. The end-user system and the client software have no knowledge about the inner workings of the system that delivers the service. The end user typically types the first URL, for example, http://www.asiaairlines.com, and then either clicks hyperlinks or completes forms on Web pages to navigate deeper into the site. In business-to-business (B2B) scenarios, the client is another server computer at the partner's site that runs an automated process and connects to exposed services on the local Internet Data Center B2B server. An example would be two servers running Microsoft BizTalk™ Server that exchange documents during the supply chain management process.

### 3.1.2 Perimeter Routers

Perimeter routers connect the infrastructure to the Internet service provider (ISP) networks. For high-end Web-business environments, full redundancy should be considered. Full redundancy requires at least two perimeter routers, with each router connected to a different ISP. This implementation provides fault tolerance and traffic aggregation. The routers should run Border Gateway Protocol (BGP) to ensure proper and fast routing. Most routers are capable of enforcing traffic policies, which should be used to help secure a perimeter network and add an additional level of security for the internal network.

### 3.1.3 Load Balancing

Load Balancing is used to distribute load among multiple servers and provide for high availability. In the Internet Data Center design, load balancing is used for the front-end Web systems and the perimeter firewalls. This design provides both resilience and scalability for the most important network elements.

### 3.1.4 Internet Facing Servers

Internet facing (front-end) servers are the collection of servers that provide the core Web services, such as HTTP and HTTPS, to Internet clients or servers. Developers usually group these front-end systems into sets of identical systems called clones. The clones run the same software and have access to the same Web content, HTML files, ASPs, scripts, and so forth, either through content replication or from a highly available file share. By load balancing the requests across a set of clones, and by detecting and separating a failed clone from the other working clones, you can achieve high degrees of scalability and availability.

### 3.1.5 Multilayer (Routing) Switches

The design can be implemented with multiple physical devices or two multilayer switches. The Internet Data Center configuration uses two multilayer switches to maintain simplicity, manageability, and flexibility. The switches are partitioned as multiple logical Layer 2 devices. VLANs are created and spanned over both switches to provide hardware fault tolerance. The servers are configured with two teamed network adapters and connected to the same VLAN on each physical switch. The traffic between VLANs is routed by using the internal router in each core switch and controlled by using access control lists (ACLs). Some network and security analysts might consider it less secure to put the external and internal networks on the same physical device. However, that would only be the case if the physical device were incorrectly configured. Most multi-layer devices are very secure, and if configured properly, do not add to the security risk. If this remains a concern, perimeter networks can simply be moved off the core switches and onto physically separate switches.

### 3.1.6 Firewalls

A firewall is a mechanism for controlling the flow of data between two parts of a network that are at different levels of trust.

### 3.1.7 Infrastructure Servers

The infrastructure VLAN was created to host domain controllers running Windows 2000 with Active Directory™ directory service and DNS. Depending on the application design, the infrastructure VLAN can also be used to host servers running components and business objects (for example, BizTalk Server 2000 or Message Queuing). If the application is designed to support three tiers, the infrastructure network can host application logic and services. Most applications are designed logically as three-tiered systems, but this design also supports a physical two-tiered application allowing Web servers to communicate directly with the servers running SQL Server.

### 3.1.8 Data and Management Servers

Back-end systems are either the actual data stores that maintain application data, or they enable connectivity to other systems that maintain data resources. Data can be stored in flat files or in database systems such as SQL Server 2000 back-end systems. The database systems are more challenging to scale and make highly available, primarily due to the data and state they must maintain.

For increased availability, a cluster supports each partition. These clusters typically consist of two nodes with access to common, replicated, or RAID-protected storage. When the service on one node fails, the other node takes over the partition and offers the service.

## 3.2 Evaluation and benchmarking criteria

### 3.2.1 Security Manageability

The current Internet Data Center design completely locks down all Web servers by using a Web server security policy and Active Directory organizational units.

Since the Web servers in the Internet Data Center design are multi-homed (two network interface cards, or NICs), architects were concerned about hackers gaining access to the production network through the back-end network interface card (NIC). The design adds another layer of protection by separating the DMZ VLAN from the rest of the production VLANs by placing a firewall directly between the internal network interface of all the servers in the DMZ VLAN and the other internal VLANs. All traffic from the DMZ VLAN that flows to the servers that are in the production VLANs must go through the firewall first. If hackers did gain access to a Web server, they would still need to beat the security configuration of the internal firewall before they could damage data.

Having a separate data and management VLAN allows the most important servers (normally the servers running SQL Server) to be placed behind two sets of protection. First, the Internet Data Center design uses stateful inspection and firewall access control lists (ACLs) to control the communication of TCP and UDP ports between servers in the DMZ VLAN and servers in the data and management VLAN. Second, the design uses VLAN technologies and additional access control lists on the switch that can be configured to control the communication of TCP and UDP ports between the infrastructure VLAN and the data and management VLAN.

### 3.2.2 Network Availability

Network availability can be achieved by providing redundancy at every level and by using automatic failover. Two network devices are implemented within each layer of the design to provide high availability at the network level. Duplicate routers, switches, and firewalls are implemented to maintain availability throughout the network. There is no single device in the design that would bring the site down. If the firewall fails, a backup firewall takes over. If one switch fails, another one takes on full load until the first one is repaired. If a Web server's network adapter fails, another NIC becomes active automatically with no impact on traffic flow. If a complete Web server fails it can be taken offline, repaired, and added back into Web farm without any impact on production. The database partitions on the SQL Server computers are protected as part of a SQL Server database cluster.

### 3.2.3 Network Scalability

Network traffic is becoming more and more unpredictable. The old 80/20 rule held that 80 percent of network traffic was limited to the workgroup, with only 20 percent involving the Internet. But with the increasing use of e-business systems, the current ratio is closer to 50/50. If trends continue the ratio may invert to 20/80, significantly increasing backbone traffic. As the Internet backbone bandwidth increases, it will increase network demand on the e-commerce sites.

Technological development is moving fast to provide technology that will ease the pressures in e-commerce networks and provide a path for upgrade to higher bandwidth requirements. The network design should include new technologies, such as Layer 2 and Layer 3 devices that switch and route traffic at wire speed. Modular and stackable switches offer port density and port speeds up to 100 megabits per second (Mbps). These devices also provide solutions for e-commerce data centers where the switch can be stacked with gigabit Ethernet (1000-Mbps) links, and provide thousands of high-speed ports.

Bandwidth aggregation for servers is available through multiple adapter technologies, which eliminate server bottlenecks by allowing incremental increases in bandwidth between a server and a switch. These technologies enable high-speed transmissions that extend the capacity of the physical medium.

# CHAPTER 4

# PLANNING AND BUILDING INTERNET DATA CENTER

## Introduction

In today's market, it is highly recommended that an Internet data center offer two different product levels: managed services and co-location services. (Cowley, 2002)

## 4.1 Managed Services

Managed services are dedicated server products built to a defined standard and offering a Common Operating Environment (COE)—standard operating system, standard network management, standard monitoring tools. Managed services are monitored and maintained in-house by the data center's own technical and support staff, with a complete maintenance and support contract. Reporting is provided to alert customers of any events and to respond to any calls for assistance from the customer.

The first step is to develop a baseline service definition for managed services. This service definition specifies the maintenance and support that are defined as standard. The definition can then be expanded to encompass "optional" or additional maintenance and support elements that can be added and charged for on an item-by-item basis.

### 4.2 Co-location Services

Co-location is the provision of racking space, power and network connectivity (frequently referred to as "power, ping and POP") to servers supplied by the customers. The attraction to the provider is that co-location offers relatively straightforward revenue generation against a minimal outlay. However, in order to be effective, co-location services must be supplied on the following basis:

- All switches and network management equipment for the co-location systems should be owned and managed by the data center
- Customers are responsible for installation and management of the equipment in the racks
- Services are governed by a clearly defined "Terms and Conditions Contract" which clearly specifies the extent to which the service is being supplied, the limitations of liability and the support and reporting from the data center to the customer.

Generally speaking, providing added monitoring services to co-location customers is not recommended. If additional services are requested, they should be negotiated on a case-by-case basis, and full consideration must be given to the additional skill sets, tools and other requirements that may be needed.

As customers supply the equipment found in co-location racks, often this equipment is poorly suited to the task and a possible risk to the data center infrastructure. An opportunity for service providers to gain additional revenue and stabilize any risk is to sell data center products to their customer, specifically products that have already been validated and deployed in other areas of the data center.

## 4.3 Advanced Services

Example: Application Service Provision (Cowley, 2002)

Front office applications are ideal candidates for the ASP deployment model. Most companies implementing front office systems need to serve a geo-graphically dispersed sales force or engineering staff and must provide reliable customer service via the Web. The operating characteristics of these applications place a premium on a reliable and centralized approach to systems management.

Just about any kind of application can be delivered by an ASP. Enabling technology from companies such as Citrix*, GraphOn* and SCO* allow current applications to be leveraged in an ASP environment. The only difference in the application (unless it was re-written for the Web) is that it is running on a central server managed by the ASP as opposed to on the end-users desktop or the company's server.

With operations managed by a service provider, companies can have the infra-structure and skills platform they need to deliver high levels of service to their distributed workforce and customers. It allows midsize companies to rapidly deploy front office applications and provides a reliable computing platform 24 hours a day, seven days a week.

For the service provider, simply hosting the application software remotely is only part of the job. The ASP has to perform a role that combines the responsibilities of an ISP, a traditional outsource service provider and a value added reseller (VAR) from which you might have purchased a non-customized software application.
In the near future, more ISPs will become ASPs; ISPs will partner with software vendors and VARs to offer ASP-applications for in-house use, rather than renting them over the Internet.

Service implementation is key. While there are minor changes to the hardware requirements for applications hosting, the ability to manage customer relations, track faults and implement change management requires specialists and solid business processes.

Customer expectations and requirements need to be defined and understood. Depending on the skills required, this can be handled by an account manager or technical project manager. This person will be the focal point of contact for all aspects of the implementation including arranging space in a racking neighborhood, requesting IP addresses and bandwidth allocation, checking availability of hardware and software and supervising the configuration and testing of equipment. When these steps are completed, this person is responsible for hand-over to the customer and the Internet data center operational environment.

## 4.4 Infrastructure: Layout

This section considers the building infrastructure requirements, such as power and lighting, that must be met in order to implement an Internet data center. (Gergg, 2001)

### 4.4.1 Building Layout

Most buildings constructed in the last ten years have been built with consideration for the requirements of computers and their support. However, in order to provide state-of-the-art, scalable Internet facilities, it is essential that any building considered for the role provides:

- Raised floors to permit adequate cabling and trunking
- Redundancy of power, such as generator systems (and possibly batteries) to support the core main supply
- Availability of fiber-optic, high-speed data connectivity

- Temperature control with separate cooling zones
- Sophisticated smoke detection and fire suppression systems
- A wide range of physical access and security safeguards (swipe card restrictions, closed circuit television monitoring, 24x7 security and security breach alarms)

To deliver the highest levels of reliability, a number of redundant subsystems are necessary. These include multiple fiber trunks coming into the building from multiple sources and multiple switching and routing of data within the building. Fully redundant power is also required on the premises, with multiple backup generators.

In addition, for a facility to be effective it is essential that it be located in very close proximity to major public and private Internet interconnects. This will keep interconnection overhead to a minimum and enable the service provider to remain competitive within the premium service marketplace.

### 4.4.2 Operation

Operating a dedicated data center environment requires a specialized team. This should include security staff to manage access to the building, as well as engineers with the skills to maintain the building infrastructure. When it comes to the network infrastructure, requirements include technical and support specialists to build and support the servers, as well as network specialists to deal with the routing, scaling and data security.

### 4.4.3 Internal Layout

Floor design and layout for housing the servers should be related to the target market sector and price of the service. Floor layout is almost always a trade-off between security, rack density, revenue potential and manageability. (Gergg, 2001)

To offer a wider choice of services to meet customer requirements, while at the same time maximizing efficiency in cabling, it is recommended that the floor layout be broken down into technical suites and racking neighborhoods. The major benefits of this approach are scaling and flexibility.
Some of the technical suites can be kept vacant and outfitted later as the data center's capacity or number of servers under management grows. And, by implementing new technical suites only when needed, the decision to equip them with racking neighborhoods, private cages or secure vaults can be deferred.
These considerations are examined more closely.

## 4.5 The Technical Suite Concept

A technical suite is an enclosed area of the data center with the infrastructure already in place to provide a secure location for hosting either managed or co-located customer systems. (Gergg, 2001)

A technical suite holds one or more racking locations and provides:

- Dedicated network trunking to all racking neighborhoods located within the suite
- Dedicated power and air-conditioning—in larger suites, AC should be zoned
- Suspended floors and ceilings for additional cable access

Lighting, fire protection and security is provided to a standard specification in each suite. Access to technical suites can be restricted via security access controls such as swipe cards.

In addition, each technical suite will normally be designated as either an area for managed or co-location systems, since each has potentially different racking layouts and power requirements. With managed services, all racking space, servers and connectivity are supplied to the customer. With co-location services, only the racking space, power and connectivity are provided. Co-location customers normally supply their own servers.

### 4.6 Secure Vaults

The use of technical suites allows the development and incorporation of secure vaults within the data center environment, if necessary. Essentially, a secure vault is a technical suite designed to provide far higher levels of client and data security than a "standard" technical suite.

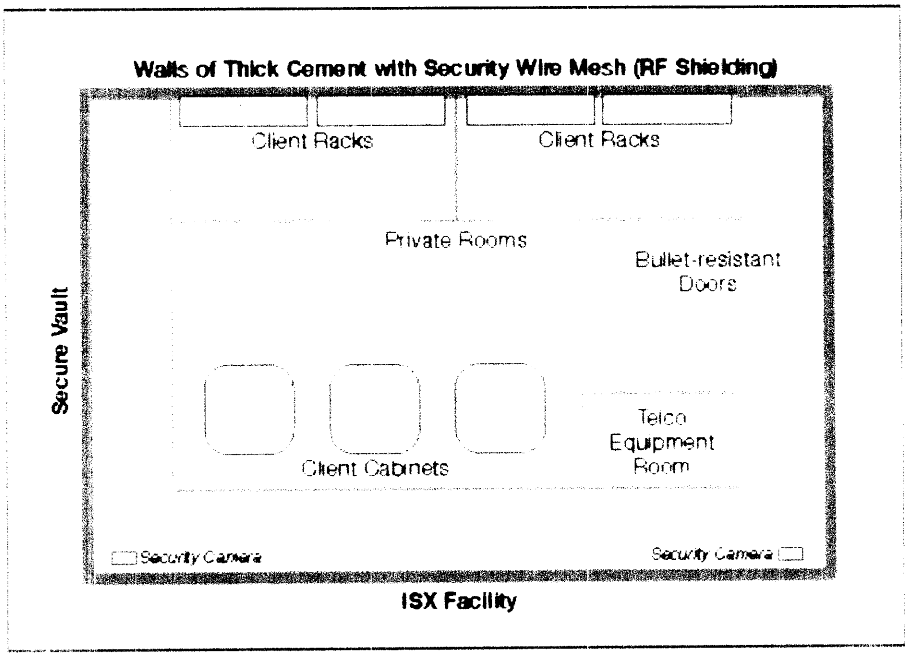

Figure 5: Typical Technical Suite "Secure vault"

*Source: Gergg (2001)*

## 4.7 Racking Neighborhoods

Racking neighborhoods are generally located within a technical suite and comprise one or more floor-mounted racks capable of supporting a number of hosting servers (Gergg, 2001). Each neighborhood within a technical suite provides:

- Dedicated network switching from the technical suite network trunking to the servers mounted in the neighborhood
- Dedicated power distribution to all racks within the neighborhood
- Localized air conditioning
- Secure, key lock access to individual racks within the neighborhood

## 4.8 Private Cages

Neighborhood racks can be optionally located in a secure private cage within a technical suite. A cage offers higher security than a standard neighborhood because access is required to the cage as well as the racks it contains. Cages are more economical for the service provider than going to the expense of setting up secure vaults.

Neighborhood racks offer the most economical use of the available space. Racking space can be rented out either as a whole or in part.

When rented as a whole, one racking unit is used per customer. This allows all of the customer's primary servers to be located together in one rack, with their backup servers located in another dedicated rack within another neighborhood. Renting neighborhood racks in part allows smaller customers to economize by paying only for the space they actually need to host their servers. This approach requires deploying mixed customers per rack.

### 4.9 Center Capacity

Total server capacity depends on the type and size of systems accommodated in each technical suite and the size of each technical suite. With space held at such a high premium in the data center, it is important to note that highest density, rack-mount systems will offer the service provider better economies. Table 2 shows typical capacity based on 10,000 square feet of floor space.

| Technical Suite | Server / 10,000 sq. ft (max) | Racks / 10,000 sq. ft (max) |
|---|---|---|
| Managed | 1500* | 400 |
| Co-location | N/A** | 400 |
| ** Based on 3.75 servers per rack – future advances in high density server design and technology should significantly increase this ratio.* | | |
| *** Dependent on the type of systems to be hosted.* | | |

Table 1: Server and Rack Capacity per Floor

*Source: Peraire (2000)*

#### 4.9.1 Facility

The building service infrastructure includes power, security, fire control and air conditioning. Based on today's market requirements, recommendations can be made in each area.

#### 4.9.2 Power Specification

Main power to an Internet data center should be supplied by the regional electric power utility. To reduce reliance on one feed, which would present a single point of failure to the business, separate feeds into the building are recommended. Once inside the building, the power should be distributed via at least two means to the individual technical suites and other protected areas. This will again prevent any single point of failure in the power supply from adversely affecting the business. One option is using two separately fed, PS-protected sets of circuits in each room and power from both at each rack.

Advances in server technology are creating more powerful and compact units that consume more power. This means the facility must be constructed to deliver more power in the future. It is strongly recommended that a minimum of 300W per square meter be provided to all technical suites, with the caveat that this requirement could be significantly higher if the density of servers exceeds that outlined in the above Table.

Power supplied to each area should be terminated in a main distribution panel. Ideally, each area has power distribution units (PDUs) within each technical suite.

These PDUs provide individual power to each rack within a suite through an individually switched and fused supply. Each rack can draw an average of 6 amps, with a maximum rating of 20 amps. Special consideration must be given to customers who have special
power requirements such as additional sockets or DC supplies.

### 4.9.3 Resilience

Higher service levels can be offered when there are a number of sources.
Competing hosting companies differentiate themselves on the quality of their power resilience. As a minimum, two-fold resilience should be provided in the form of:

- Uninterruptible power supply (UPS) to each area or neighborhood. This includes a recommended minimum (or equivalent) of 600 Kva capacity supplying 3 phase 240v AC to each neighborhood and a UPS battery run time of several minutes. The longer the run time, the greater confidence customers will have in power availability. Run time should at least equal the time it takes to bring the generators on line.

- Diesel generator backup configured to start automatically within seconds after a main power source failure to provide power to all relevant services. The generator should be running at full load within the UPS battery run time. Diesel fuel must be readily available on-site for continuous generator running over several hours and available off-site for continuous running over several days as a contingency. On-site storage tanks should be capable of being filled while the generators are running.

### 4.10 Building Security and Access Control

Building and network security is a highly visible component of service and very important to customers. With that in mind, a robust security policy supported by a workable set of procedures, skills and tools is essential.

The security policy is required to manage access to the data center and monitor activity within the building. This activity is normally needed 24 hours a day, every day of the year. This should be supported, where required, by the use of closed-circuit television (CCTV) to monitor the exterior of the data center and also the corridors and technical suites within the building. Physical access to technical suites and other areas of the building should be controlled and monitored on an ongoing basis, preferably by a swipe card facility to maintain and control access to restricted areas.

All staff joining the data center should be security screened and subject to the same access controls as visitors to the building—including swipe card access.

### 4.10.1 Fire Control

It is crucial that the building be protected by a fully automated fire detection and suppression system. These systems are typically zoned throughout the building and linked into a battery backup system. An ideal system would provide automated detection, announcement and control of a fire condition before damage occurs. A manual override to the system is recommended.

All technical suites should be equipped with FM200 (or equivalent) gaseous extinguishing systems. These systems are designed to provide rapid discharge and flame suppression in the event of a fire. This minimizes the damage to equipment and reduces danger to personnel. The chemical deploys after a 30-second countdown and slows the fire by preventing combustion.

### 4.10.2 Air Conditioning

Equipment performance and life span can be significantly improved by housing the system under optimum environmental conditions. Typically, this should be around a constant 68° F, plus or minus 3° F, with humidity at a constant 45%-50%. Cooling units should be placed over walkways or hallways, never over racks.

### 4.10.3 Staff Facilities

Considerations need to be given to the provision of staff facilities. Many of the personnel working at the data center will be required to work on shift, be on call or put in longer than average hours. It is therefore recommended that staff should be provided with a rest area away from the main technical suites, including a kitchen facility. Overnight secure parking is recommended for on-call staff who may not be able to rely on public transportation.

## 4.11 Systems Infrastructure

Designing network access into the data center requires significant commercial and competitive consideration—greater resilience is gained by having multiple carriers (Telco & ISP). The cost of band-width is a major driver for investing in products that allow more efficient management of bandwidth in the facility. Products such as caching technologies, load balancers and switches should be reviewed and deployed.

Typically, fiber cable running over multiple SDH rings and dark fiber ducts enters the data center at both ends of the building. Ideally, the cables will come from competing Telcos and different Telco exchanges. It is then distributed throughout the building using data risers at opposite ends of the structure. This enhances the resilience of the telecom network infrastructure. Data access should also offer multiple connections to the Internet—such as nodes on the ISP's network.

Within the data center, the network infrastructure should be based on Gigabit topology, which offers scalability, resilience and manageability. Another proven technology, ethernet is economical, highly scalable (10/100/1000 MB) and has a large user base. This means there are many vendor solutions avail-able, assuring competitive pricing and good quality equipment.

## 4.12 Rack Configurations

Within a technical suite, servers are located within racking neighborhoods. Each server within a rack is supplied with power and network connectivity. In addition, managed servers are provided with monitoring, management and backup facilities.

A neighborhood comprises a number of server racks supported by dedicated switch racks for network connectivity and a pre-defined number of network connections per server (for managed services) or per customer (for co-location services). The exact configuration for a neighborhood is dependent on the equipment being used. It is important to note that highest-density rack-mount systems designed for the data center will offer the service provider better economies than traditional, general-purpose servers.

When evaluating server platforms, it is absolutely essential that they provide seamless interoperability with the operating systems, development tools and applications needed to run a successful data center

## 4.13 Data Center Management and Operation

### 4.13.1 Service Management Center (SMC)

The Service Management Center is the core of the data center facility, providing systems management for all managed services and monitoring for the network. Correctly set up and managed, the SMC provides first-level support (1LS) for all alerts, incidents and problems, first-level contact for customers with the data center, direct feedback to customers on incident and problem resolution and dedicated network management and monitoring. (Peraire, 2000)

Peraire (2000) quoted that, the SMC should be located within the data center itself and staffed 24 hours a day, seven days a week. During core hours, support should be provided as follows:

- SMC staff: incident and problem logging, first-level support/resolution
- Technical support staff: second-and third-level support

Outside of core support hours, the SMC staff should still provide initial logging and first-level support, with technical support staff providing second- and third-level support on an on-call basis.

In order to be effective, the SMC must be equipped with the following:

- Dedicated management and monitoring tools for all operational managed services (e.g. HP OpenView*, Utracomp UltraFrame Works* or equivalent)
- Automated backup facilities for all managed services
- An on-line Call Management System (e.g. Vantive*, Remedy*, Quetzal*)
- An integrated Change Management/ Asset Management/Configuration Tool (e.g., Utracomp Red Box*)
- If possible, integrated Call Management/ Help Desk Management with Change Control (e.g., via Ultracomp Red Box or via a process management tool such as InformAction's Matrix* product)

### 4.13.2 Service Monitoring and Maintenance

The major selling point for dedicated managed services is the ability to provide "package" standard service monitoring facilities and to offer customers more proactive monitoring and auto-correction of faults. This requires the deployment of the appropriate monitoring tools and software. (Peraire, 2000)

Basic monitoring allows the SMC to check whether the network or a server is up or down, to check the general health of the network in terms of parameters such as packet loss and to view all log files. Basic monitoring could also cover server metrics such as hard disk usage, CPU and memory usage. Incidents detected by this basic monitoring would then be entered into the Call Management System, allowing them to be resolved via the defined incident and problem management process.

Beyond the basic level, proactive monitoring and maintenance provides "value added" services to the managed service environment. Examples of proactive service include trends monitoring, auto-mated responses to given conditions, new patches/service packs firewall monitoring for intrusions and input to the call logging system when required.

These services can be provided on a chargeable basis, with customers paying a monthly fee to have the tasks carried out over and above the basic monitoring package. Or, customers could pay for them on an ad-hoc, fixed-price basis.

### 4.13.3 Customer System Backups

Backups should only be provided for Managed Services. (Peraire, 2000) Furthermore, in order to be effective, it is essential that any backup solution deployed is capable of backing-up files and tables that may be open at the time the backup is made. To this end, it is recommended that a product such as OpenFile* be evaluated as a part of the backup solution. (Peraire, 2000)

#### 4.13.3.1 Storage Area Networks

The typical computing environment, in which the backup media is provided on a per-server basis, lacks the flexibility and scalability required when transparently backing up vast quantities of data for a cross-section of clients. Therefore, alternative methods of secure, transparent backup need to be implemented.

A relatively new approach to backing-up data, which offers significant advantages to the data center environment, is the Storage Area Network (SAN) solution.
SANs are based on a network of fiber channels to facilitate high speed and switches connecting storage devices (i.e., disk arrays, optical disks, tape libraries) to servers on a many to many basis. In an Internet hosting environment, this requires the server to have two ports: a public port for Internet access and a private port for backups and management.

According to (Peraire, 2000), SANs have a number of advantages:

- Facilitates universal access and sharing of resources
- Supports unpredictable, explosive information technology (IT) growth
- Provides affordable 24x365 availability
- Simplifies and centralizes resource management
- Improves information protection and disaster tolerance
- Enhances security and data integrity of new computing architectures

#### 4.13.3.2 SAN and Network Attached Storage

SAN and Network Attached Storage (NAS) provide similar facilities. The key differences can be summarized as:

- Storage Area Networks enable multiple servers to share central Fibre Channel RAID storage for higher performance, lower management cost and provide unlimited capacity growth. As stated above, SANs usually have dedicated network connectivity
- Network Attached Storage provides direct ethernet attachment of RAID storage without any disruption or downtime to existing servers

Problem Management, Configuration Management and Change Control

Problem, configuration and change management are vital to the successful implementation of any managed service environment.

- Problem management, through the implementation of a Call Management System (CMS), enables the Service Management Center to log, track and resolve incidents either as they occur or as customers report them

- Through the implementation of a Configuration Management Database (CMDB), an Internet data center can baseline build configurations (hardware, operating system and software) of all managed services. Updates and changes to individual configurations can be tracked
- Change Management enables correct logging and implementation of hardware and software upgrades, changes to the operational parameters in the data center and changes to monitoring services. This ensures that there is minimal impact on current services or to customers

Problem, configuration and change management are closely interrelated.
Because of the way they are inter-related, and given the fact that correctly implemented problem and configuration management are a major plus in selling managed services, it is recommended that:

- Adequate planning should be given to the provision of all three (problem management, configuration management, change management)
- All three should be linked as closely as possible via their supporting software, the CMS and the CMDB
- All three should be implemented to standards such as those set out in the CCTA IT Infrastructure Library (ITIL)

## 4.14 Data Center Organization

### 4.14.1 Structure

Data center management can be divided into two key areas, production and development.

Production is the actual provision of services to customers from the initial sales contact through the implementation and monitoring of services. The production staff is responsible for:

- Managing the pre-sales process to ensure customers receive a service they require
- Managing current product portfolio
- Advising the development staff on emerging requirements from customers
- Building servers to current configuration standards
- Installation and implementation of new customer services
- Managing and maintenance of implemented services
- Providing first-level problem management
- Providing customer support

Development is defined as the ongoing tactical and strategic development of products and infrastructure to meet the demands of an evolving marketplace.
The development staff is responsible for:

- Defining product baseline build configurations
- Product development in conjunction with the pre-sales process of the production staff
- Deployment and integration of new products
- Defining, testing and deploying tool sets for use by the SMC and support staff
- Network infrastructure management
- Management strategy definitions including backup strategies and disaster recovery strategies
- Providing second- and third-level problem management

Production and development areas are closely interlinked by a number of processes. One of these processes is problem management. Production staff provides initial support such as logging calls and giving customer feed-back, plus first-level support through the Service Management Center (SMC). Development staff provides second- and third-level support via a call escalation process for those problems that cannot be resolved by the SMC.

Other processes where the two areas are interlinked include change management and configuration management. Development staff defines the baseline configuration standards for server hardware, operating system builds, software toolkits, network infrastructure and connectivity.

This information is entered into the Configuration Management Database, and individual updates to specific configurations are then made to the CMDB as Production support and maintain services in the "live" environment. Change management is the control tool to ensure all configuration changes are trapped, recorded and implemented.

Still another area is support and integration. Development staff defines the tactical and strategic development of data center tools and then provides the specialization and support needed to implement new technologies into the production environment, working
in close co-operation with the Service Management Center.

### 4.14.2 The Production Environment

The production environment is sub-divided into three areas: pre-sales, project management and the Service Management Center. (Peraire, 2000)

Pre-sales is a vital function that provides a direct link between the sales force responsible for obtaining new customers and the technical staff responsible for providing services to those customers. Frequently, potential customers request services that are beyond the standard managed service provided by a data center. The role of the pre-sales area is to capture such requests from the sales force and ensure adequate technical input is given to a potential customer's requirements before any sale agreement is made.

According to (Gergg, 2001), there are several benefits of this kind of approach:

- Customers are not "sold" solutions that cannot be easily met by the data center
- The data center is able to judge customer requirements on a case-by-case basis and assess whether it is cost-effective to update the standard managed services to include frequently requested items
- The sales force is not required to undertake a technical evaluation of a potential customer's specific requirements—instead, the sales force can request assistance from technical pre-sales staff

While the development environment has overall responsibility for defining the services supplied by the data center, implementing those services on behalf of customers is a specialized task requiring the services of a dedicated implementation team of engineers. This implementation or project management team has the responsibility of:

- Building services to the required specification
- Updating the CMDB as new builds are completed
- Implementing new services for customers
- Handing over new services to the Service Management Center

The most efficient way to implement services for new and existing customers is to handle them on a project basis. Each service or customer is classified as a distinct project overseen by a technical project manager. This project manager stays in contact with the  customer during the build and implementation process, allowing the implementation engineers to concentrate on building, testing and implementing the required servers and services.

The pre-sales and project implementation processes govern all work to implement customers and services within the production environment. These processes, which should be developed to a defined standard, specify all the required steps that must be followed in order to successfully sign up a customer with the data center and then implement the services the customer has purchased. There should be a logical flow of information from the pre-sales environment, through to the hand-over of a service to the Service Management Center.

Hand-over Documentation for each customer is one of the final steps in the implementation process. This document should include:

- Support and contact information
- Service description
- Hardware and software configuration
- Password and access control information
- Standard service monitoring and maintenance (managed services only)
- Incident and problem reporting procedures (managed services only)
- Backup methodology (managed services)
- New software installation guidelines (managed services only)

The Service Management Center must be staffed 24 hours a day, seven days a week throughout the year. Staff will therefore be required to work a shift pattern, and a suitable shift system must be established. Since the SMC will provide initial incident investigation and resolution, the staff employed need to have a thorough understanding of the systems and services supplied by the data center. Preferably, they should also be familiar with the customers they will be dealing with.

Implementation and support staff will not be required to work shifts, but will be required to provide second- and third-level support. Therefore, staff employed in these areas will be required to work on an on-call rotation.

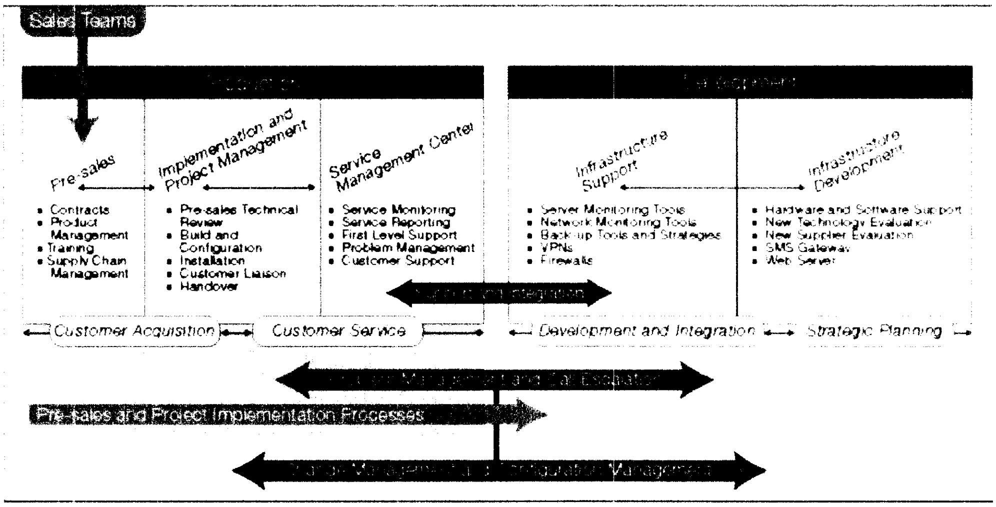
Sales Teams
Pre-sales
Implementation and Project Management
Service Management Center
Infrastructure Support
Infrastructure Development
• Contracts
• Product Management
• Training
• Supply Chain Management
• Pre-sales Technical Review
• Build and Configuration
• Installation
• Customer Liaison
• Handover
• Service Monitoring
• Service Reporting
• First Level Support
• Problem Management
• Customer Support
• Server Monitoring Tools
• Network Monitoring Tools
• Back-up Tools and Strategies
• VPNs
• Firewalls
• Hardware and Software Support
• New Technology Evaluation
• New Supplier Evaluation
• SMS Gateway
• Web Server
Customer Acquisition
Customer Service
Development and Integration
Strategic Planning
Pre-sales and Project Implementation Processes

Figure 6: Internet Data Center Organization

*`Source: Gergg, 2001*

### 4.14.3 The Development Environment

The development environment is sub-divided into two areas: infrastructure support and infrastructure development.

Infrastructure Support Staff is responsible for integrating products and technology into the current production environment. This staff also has responsibility for providing second-level support to the Service Management Center to resolve incidents and problems via the problem management process.

Infrastructure Development is the area that evaluates emerging products and technology and reviews new standards for their benefit to the data center. This area also provides the highest level of support, third-level, for all hardware and software implemented within the data center. Technologies evaluated and approved for use within the data center are passed to the infrastructure support group for integration into the data center, with support from the development team.

## Conclusion

If customers are to outsource their mission-critical Internet operations, they expect a physical and technical environment that offers the highest levels of reliability and flexibility. Accordingly, the Internet data center must provide the physical environment,
network connectivity, technical skills and server hardware necessary to keep Internet servers up and running 24 hours a day, seven days a week.

Establishing and operating an Internet data center requires a high degree of planning and implementation, not just to ensure the integrity of the data center environment itself, but also in the definition and integration of the supporting infrastructure—sales teams, sales
support, customer support, etc.

One key to establishing a successful data center must be that of scalability.
Facilities do not, in the first instance need to be large—but they do need to offer the potential for rapid growth.

# CHAPTER 5

# FUTURE DIRECTION AND FINDINGS

**Introduction**

The key architectural elements of the Internet Data Center network include clients, perimeter routers, load balancers, cloned front-end Web servers, multilayer switches, firewalls, infrastructure servers and back-end database and management systems. (Microsoft, 2002) This chapter focuses on the logical components that provide an infrastructure that is scalable, available, secure, and manageable.

## 5.1 Architectural Elements

### 5.1.1 Internet Clients

In the Internet Data Center environment, clients issue requests to a service name, which represents the application being delivered to the client. The end-user system and the client software have no knowledge about the inner workings of the system that delivers the service. The end user typically types the first URL, for example, http://www.blueyonderairlines.com, and then either clicks hyperlinks or completes forms on Web pages to navigate deeper into the site. In business-to-business (B2B) scenarios, the client is another server computer at the partner's site that runs an automated process and connects to exposed services on the local Internet Data Center B2B server. An example would be two servers running Microsoft BizTalk™ Server that exchange documents during the supply chain management process.

### 5.1.2 Perimeter Routers

Perimeter routers connect the infrastructure to the Internet service provider (ISP) networks. For high-end Web-business environments, full redundancy should be considered. Full redundancy requires at least two perimeter routers, with each router connected to a different ISP. This implementation provides fault tolerance and traffic aggregation. The routers should run Border Gateway Protocol (BGP) to ensure proper and fast routing. Most routers are capable of enforcing traffic policies, which should be used to help secure a perimeter network and add an additional level of security for the internal network.

### 5.1.3 Load Balancing

Load Balancing is used to distribute load among multiple servers and provide for high availability. In the Internet Data Center design, load balancing is used for the front-end Web systems and the perimeter firewalls. This design provides both resilience and scalability for the most important network elements.

### 5.1.4 Internet Facing Servers

Internet facing (front-end) servers are the collection of servers that provide the core Web services, such as HTTP and HTTPS, to Internet clients or servers. Developers usually group these front-end systems into sets of identical systems called clones. The clones run the same software and have access to the same Web content, HTML files, ASPs, scripts, and so forth, either through content replication or from a highly available file share. By load balancing the requests across a set of clones, and by detecting and separating a failed clone from the other working clones, you can achieve high degrees of scalability and availability.

### 5.1.5 Multilayer (Routing) Switches

The design can be implemented with multiple physical devices or two multilayer switches. The Internet Data Center configuration uses two multilayer switches to maintain simplicity, manageability, and flexibility. The switches are partitioned as multiple logical Layer 2 devices. VLANs are created and spanned over both switches to provide hardware fault tolerance. The servers are configured with two teamed network adapters and connected to the same VLAN on each physical switch. The traffic between VLANs is routed by using the internal router in each core switch and controlled by using access control lists (ACLs). Some network and security analysts might consider it less secure to put the external and internal networks on the same physical device. However, that would only be the case if the physical device were incorrectly configured. Most multi-layer devices are very secure, and if configured properly, do not add to the security risk. If this remains a concern, perimeter networks can simply be moved off the core switches and onto physically separate switches.

### 5.1.6 Firewalls

A firewall is a mechanism for controlling the flow of data between two parts of a network that are at different levels of trust.

### 5.1.7 Infrastructure Servers

The infrastructure VLAN was created to host domain controllers running Windows 2000 with Active Directory™ directory service and DNS. Depending on the application design, the infrastructure VLAN can also be used to host servers running components and business objects (for example, BizTalk Server 2000 or Message Queuing). If the application is designed to support three tiers, the infrastructure network can host application logic and services. Most applications are designed logically as three-tiered systems, but this design also supports a physical two-tiered application allowing Web servers to communicate directly with the servers running SQL Server.

### 5.1.8 Data and Management Servers

Back-end systems are either the actual data stores that maintain application data, or they enable connectivity to other systems that maintain data resources. Data can be stored in flat files or in database systems such as SQL Server 2000 back-end systems. The database systems are more challenging to scale and make highly available, primarily due to the data and state they must maintain.

For increased availability, a cluster supports each partition. These clusters typically consist of two nodes with access to common, replicated, or RAID-protected storage. When the service on one node fails, the other node takes over the partition and offers the service.

### 5.1.9 Corporate Connection

The corporate connection is the link between the Internet Data Center and the internal network used for the enterprise.

## 5.2 Network Design and Goals

The Internet Data Center design is flexible because it uses VLAN technologies to separate servers and communication traffic. The core VLANs created in support of the different server traffic requirements consist of DMZ, infrastructure, and data/management VLANs. (Microsoft, 2002)

The DMZ, infrastructure, and data and management VLANs are the core VLANs for the design decisions made in the architecture. (Microsoft, 2002) The DMZ network consists of multi-homed Web servers and the external DNS servers that users can query directly from the Internet through the Internet perimeter firewalls. The DMZ is in fact made up of three VLANs, to allow for traffic isolation, but for now we will simply refer to all three as the DMZ VLAN. The data and management VLAN is comprised of the SQL Server 2000 database servers and other required management and backup servers. The infrastructure VLAN contains servers that provide services required by the DMZ network or data and management VLANs, for example, Active Directory and DNS. Additional services will be added here as the application for the Internet Data Center is developed.

The following diagram (figure 7) depicts the network segmentation and inter-communication in the form of a logical design diagram.

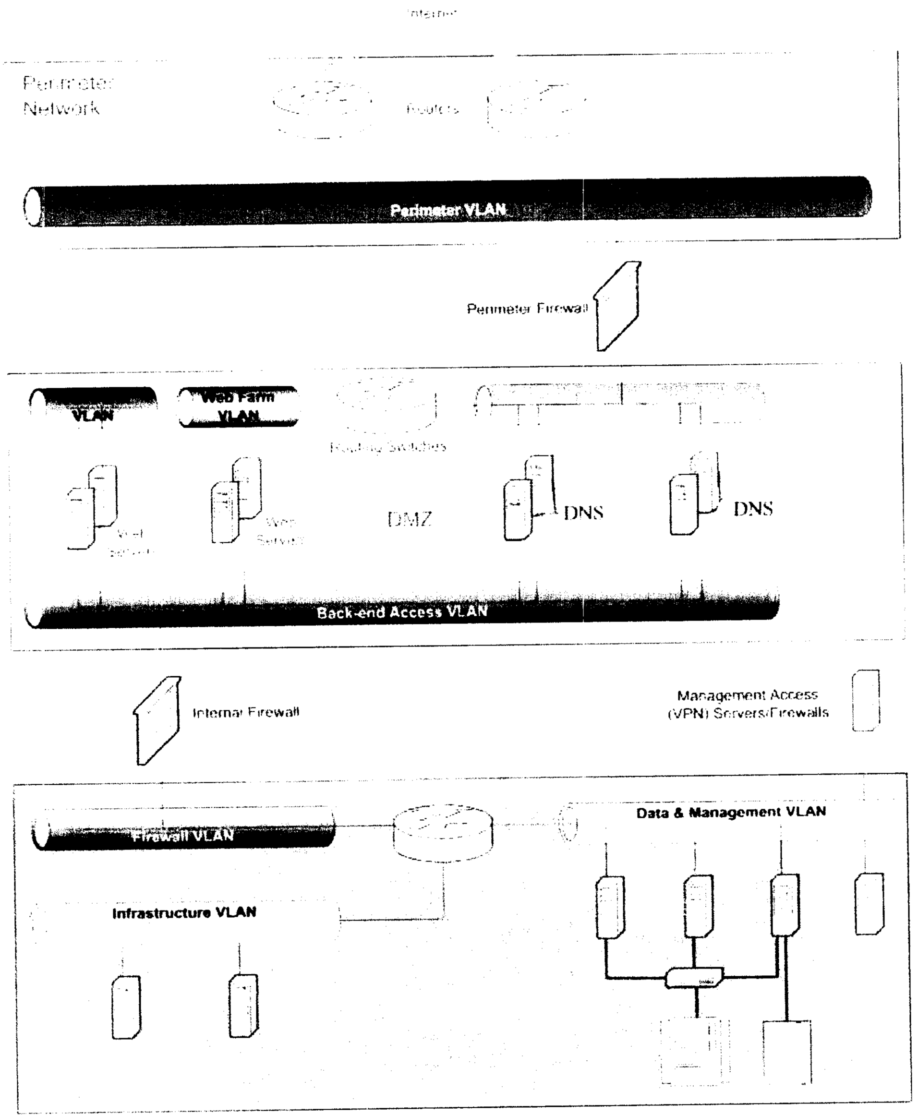

Figure 7: Internet Data Center Conceptual Model

Before going into the details of each of the components within the VLANs, it is useful to consider why the design looks the way it does. The following sections discuss some aspects of the design.

### 5.2.1 Traffic Flow Manageability

The VLAN designallows traffic flow to be managed efficiently by creating a series of protected security devices to which rules and policies can be applied.

### 5.2.2 Security Manageability

The current Internet Data Center designcompletely locks down all Web servers by using a Web server security policy and Active Directory organizational units.

Since the Web servers in the Internet Data Center designare multi-homed (two network interface cards, or NICs), architects were concerned about hackers gaining access to the production network through the back-end network interface card (NIC). The designdesign adds another layer of protection by separating the DMZ VLAN from the rest of the production VLANs by placing a firewall directly between the internal network interface of all the servers in the DMZ VLAN and the other internal VLANs. All traffic from the DMZ VLAN that flows to the servers that are in the production VLANs must go through the firewall first. If hackers did gain access to a Web server, they would still need to beat the security configuration of the internal firewall before they could damage data.

Having a separate data and management VLAN allows the most important servers (normally the servers running SQL Server) to be placed behind two sets of protection. First, the Internet Data Center design uses stateful inspection and firewall access control lists (ACLs) to control the communication of TCP and UDP ports between servers in the DMZ VLAN and servers in the data and management VLAN. Second, the design uses VLAN technologies and additional access control lists on the switch that can be configured to control the communication of TCP and

UDP ports between the infrastructure VLAN and the data and management VLAN.

### 5.2.3 Network Availability

Network availability can be achieved by providing redundancy at every level and by using automatic failover. Two network devices are implemented within each layer of the designto provide high availability at the network level. Duplicate routers, switches, and firewalls are implemented to maintain availability throughout the network. There is no single device in the design that would bring the site down. If the firewall fails, a backup firewall takes over. If one switch fails, another one takes on full load until the first one is repaired. If a Web server's network adapter fails, another NIC becomes active automatically with no impact on traffic flow. If a complete Web server fails it can be taken offline, repaired, and added back into Web farm without any impact on production. The database partitions on the SQL Server computers are protected as part of a SQL Server database cluster.

### 5.2.4 Network Scalability

Network traffic is becoming more and more unpredictable. The old 80/20 rule held that 80 percent of network traffic was limited to the workgroup, with only 20 percent involving the Internet. But with the increasing use of e-business systems, the current ratio is closer to 50/50. If trends continue the ratio may invert to 20/80, significantly increasing backbone traffic. As the Internet backbone bandwidth increases, it will increase network demand on the e-commerce sites.

Technological development is moving fast to provide technology that will ease the pressures in e-commerce networks and provide a path for upgrade to higher bandwidth requirements. The network design should include new technologies, such as Layer 2 and Layer 3 devices that switch and route traffic at wire speed. Modular and stackable switches offer port density and port speeds up to 100 megabits per second (Mbps). These devices also provide solutions for e-commerce data centers where the switch can be stacked with gigabit Ethernet (1000-Mbps) links, and provide thousands of high-speed ports.

Bandwidth aggregation for servers is available through multiple adapter technologies, which eliminate server bottlenecks by allowing incremental increases in bandwidth between a server and a switch. These technologies enable high-speed transmissions that extend the capacity of the physical medium.

### 5.2.5 Simple Architecture

To simplify the Internet Data Center design, all unnecessary VLANs have been eliminated and multi-homing is used only when absolutely necessary. In particular, the design has no separate management VLAN because this would require that both the infrastructure, management and the data servers be multi-homed.

By placing all of the management servers where they have the most impact, the Internet Data Center design eliminates the complexity and addresses some of the traffic and security concerns that having a separate management network would introduce. However, the design does have multi-homed Web servers, because it is important to segment inside and outside traffic with separate VLANs and eliminate NetBIOS over TCP/IP on the outside interface. In other words, the Internet Data Center design does incorporate complexity when it is justified by a significant gain in availability, reliability, management, or security.

## 5.3 Router Design

The connection point between the Internet Data Center network and the outside world is the router. These perimeter routers (also known as border or edge routers) must enable the main services of any network design: security, high availability, and scalability.

### 5.3.1 Internet Perimeter Router

In the Internet Data Center design, the perimeter routers operating system security services provide the first step to a secure front end to the network. This is achieved by using the extended ACLs of the routers to secure the network traffic allowed onto the perimeter network.

For reliability and availability, the network uses a high availability protocol to ensure that this router configuration is fault tolerant. The Border Gateway Protocol (BGP) provides routing availability and load-balancing capability. The perimeter routers also provide a set of QoS (Quality of Service) features that could be used to improve the availability of user sessions during times of peak load on the network.

The Internet Data Center design uses perimeter routing to:

- **Implement redundant routers** for the Internet Data Center design to eliminate the single point of failure. Connect each router to a different ISP connection for maximum availability.

- **Supply BGP capability** to fully use ISP routing information. This is critical in a multiple ISP scenario, where network load balancing and policing routing is important. In addition, routers with BGP capability are recommended for scalability.

  **Note**: This requires the ability to obtain an Autonomous System Number (ASN)

- **Create multiple paths through the network infrastructure** for higher availability and make use of these paths to allow for load sharing and higher scalability through routing-protocol load balancing.

- **Use external BGP routes** (EBGP on perimeter routers) for the propagation of local IP network routes to the interconnected ISPs. This allows path discovery to the e-commerce site. By exchanging full Internet BGP routes with all ISPs, the perimeter routers can determine the best return path and offer the quickest response to the customer.

- **Apply tight extended ACLs** from the inbound interfaces to the perimeter routers. These ACLs should only allow traffic that is relevant to the e-commerce site.

- **Deny any traffic destined to the routers by using ACLs**, but allow BGP traffic that uses TCP/179 if the packets arrive from adjacent ISP routers.

- **Prevent Internet Control Message Protocol (ICMP)** packets from being transmitted through the router, because support of the ping command, and similar capabilities, can lead to attacks.

- **Install a spoofing ACL** to prevent traffic that is structured to appear as if it were sourced from the data center from actually entering the data center. This ensures that traffic with a source address in the perimeter network really does originate there and not from the outside.

- **Secure the console interface** on the routers themselves with user names and passwords. One method is to use Remote Access Dial-In User Service (RADIUS), to authenticate and account for those administrators who log on to the router consoles. Use Kerberos authentication or Secure Shell (SSH) to access the router console.

- **Allow only TCP/80 (HTTP), TCP/443 (SSL), TCP/25, and UDP/53 (DNS) to enter the data center**. It is possible that customized applications that were developed to run on top of the Internet Data Center design will require additional protocols to allow the clients to perform additional actions, such as using File Transfer Protocol (FTP).

### 5.3.2 VPN Perimeter Router

The second router configuration in the Internet Data Center design was created to provide a secure connection to the design from remote locations. These connections are primarily for access to the management systems for the support staff. This allows for the separation of the management traffic from the production traffic so that any management tasks cannot impact the customer's bandwidth to the site. However, to simplify the configuration and to cut down on the hardware requirements of the system, the design of this connection could be incorporated into the existing Perimeter Router design.

## 5.4 DNS Design

The DNS service design is a very important part of the Internet Data Center design. The base Internet Data Center design implements Windows 2000 DNS services in a design commonly known as “Split DNS”. The split DNS design consists of external DNS servers that provide name resolution to the Internet clients and internal DNS servers that service only the internal name space. The following diagram shows of how the internal and external DNS namespaces are split.

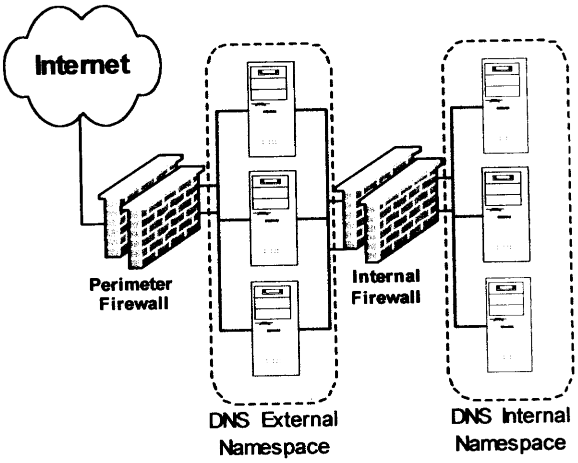

Figure 8: Split DNS configuration

### 5.4.1 External DNS Services

Zone transfers are configured to take place between the primary and secondary external DNS servers on the internal side of the multi-homed servers. This prevents any exposure of zone data on the Internet-facing side of the servers. Zone transfers are not allowed between the internal and external DNS servers. This separation or split DNS configuration allows internal DNS namespace information to remain isolated and unavailable to the Internet. Under no circumstances should DNS queries from the Internet be allowed to pass through the DMZ into the internal network for name resolution. Zone transfers are also not allowed between the external DNS servers and any DNS server on the Internet. In certain limited cases, you may want to transfer your external DNS zone to your ISP DNS server for additional redundancy. If implemented, the DNS servers should be configured to only permit zone transfers between the ISP DNS server and your external DNS servers.

### 5.4.2 Internal DNS Services

Internal sets of DNS servers, located on the infrastructure VLAN, are used to support Active Directory and support name resolution for the servers in the DMZ infrastructure, and data and management VLANs.

For purposes of scalability and redundancy, multiple DNS servers are configured in the architecture. For economic and simplicity reasons, the DNS servers are installed on the Active Directory domain controllers and are configured to integrate their zone files with the Active Directory database. By configuring the DNS servers as Active Directory integrated, both DNS servers' zone files will automatically be replicated by Active Directory, and both DNS servers will be able to manage the same namespace by acting as primary DNS servers. All domain member servers are configured to point to both DNS servers as primary and secondary. All servers can register their friendly names to the DNS servers automatically through Dynamic DNS (DDNS). In most configurations, including the Internet Data Center architecture, internal systems must be able to communicate to the systems outside of the intranet. To accomplish Internet name resolution, internal DNS services are configured to forward all DNS queries to the external DNS servers for resolution.

### 5.4.3 Split-Split DNS Design

For additional security, a split-split DNS design can be implemented. Split-split DNS takes a split-DNS design one step further by separating the Advertiser and Resolver services in the external DNS. Advertiser services in DNS handle queries from clients on the Internet for zones that the DNS server is authoritative for. By limiting queries to local zones and disabling recursion, these services are protected from attacks. Resolver services in DNS handle queries forwarded from internal DNS servers and resolve the requests on behalf of the internal servers. Resolver servers do not listen for queries from the Internet. They only listen for requests from the internal DNS servers. In this way, they are protected from attacks originating from the Internet.

Cache poisoning is one type of attack that a split-split DNS design guards against. DNS poisoning involves providing bogus data to a DNS server in order to misdirect users. Cache poisoning works by sending a DNS query to a DNS server on the Internet for a zone that the DNS server is not authoritative for. If recursion is enabled on the DNS server, it will attempt to resolve the request for the client by finding the DNS server on the Internet that is authoritative for the zone the client is requesting data from. The DNS server then sends a request to this authoritative DNS server to resolve the request for the client. Unfortunately, this authoritative DNS server is owned by a malicious individual or organization, which in its response sends a valid answer, but also sends an attack on the end of the answer that could potentially alter the DNS servers cache. In a split-split DNS design, all servers listening on the Internet have recursion disabled, so this type of attack is prevented.

## 5.5 Load Balancing Design

Load Balancing is a way of distributing network traffic amongst servers in a group to provide scalability and availability. IDC heavily leverages load balancing strategy to provide the increased availability and scalability of many architectural elements in the design.

### 5.5.1 Load Balancing Mechanisms

Load Balancing is typically achieved through one of three methods (or some combination of the three). These methods and their trade-offs are captured in the following table:

| **Method** | **Advantages** | **Disadvantages** |
|---|---|---|
| Round-robin Domain Name System (DNS) | • Simple<br>• Cheap (most DNS servers allow this functionality) | If a server fails, the host needs to be manually removed from the DNS. The failed address could also be cached in DNS servers all over the Internet causing clients to be sent to a failed server. This situation is not resolved until the node is either brought back online, or DNS server caches expire the failed entry and perform a new lookup. |
| Hardware-based network load balancing | • Fast, intelligent (will dynamically remove failed nodes from the cluster)<br>• Optimizes network utilization<br>• Advanced load-balancing – content switching | • Expensive<br>• Becomes a single point of failure unless a redundant hardware is used, further adding to the cost |
| Software-based Network Load Balancing | • Fast<br>• Offers some native intelligence. For example, a failed node will automatically be removed from the cluster if connectivity to the node is lost.<br>• No single point of failure (the distributed nature of | • No native detection of application-level failures<br>• Relies upon network "flooding" because every node in the cluster must receive every packet<br>• Clusters cannot span subnets<br>• Recommended cluster sizes are 12 nodes (although Network Load Balancing will allow the configuration of up to 32 nodes per |

| Method | Advantages | Disadvantages |
|---|---|---|
| | load balancing eliminates this)<br>• Included with Windows 2000 Advanced Server, Windows 2000 Datacenter Server, and Application Center 2000 | cluster, Application Center 2000 will allow no more than 12 nodes in a cluster). Using 12 node clusters, and then using round-robin DNS to load-balance between cluster IPs improves scalability.<br>• If ISA Server is the perimeter firewall then round-robin DNS is not required. ISA server will manage distributing traffic to multiple load balanced IIS cluster farms using one external IP address. |

*(Source: Microsoft, 2002)* Table 1: Load balancing mechanisms

### 5.5.2 State Management

Since HTTP is a stateless protocol, web developers have always had to figure out a way to address this limitation in order to build more functional applications. The realization of this problem has prompted web server vendors to add mechanisms to deal with this issue. IIS, for example, provides session variables that can be used very effectively to maintain state. However, this approach has disadvantages when load balancing is introduced. Session variables for a specific session will only exist on a given server. Should the particular user get redirected to a different server, the user's session information is lost. In order to address this, the concept of affinity was introduced. Affinity, as its name implies, allows the user session to get assigned to a particular server. Affinity permits session information to remain available to application as all requests from a specific user are directed to the specific web server. Affinity is usually based on connection information (source and destination IP addresses and ports). This is usually very effective. However, some larger ISPs may internally reroute client traffic so the same user may appear to come from a different proxy server (different source IP address) or a different proxy server array (different source IP subnet), thus potentially breaking affinity. The recommended way of maintaining state in a SQL database or a file share on the back end can mitigate all of these issues. However, should your application not leverage this strategy, you can still load balance by using either a request proxy or a content switch. A request proxy allows clients to make a request against any server in the web farm, even if the user has not communicated with the particular web server. The web server checks a cookie that indicates which server has state information, proxies the request to the target server on user's behalf, and returns information to the user. This functionality is found in the Application Center 2000 product. Note that both request proxy and content switching are still less optimal than backend state management since a failure of a particular server containing state will result in loss of that session state.

### 5.5.3 Content Switching

Content switching is yet another way to mitigate the state management problem. Content switching is a more intelligent way of load balancing the request by allowing information from layers 4-7 of the OSI model contained in the request itself (requested URL and cookies) to be used to direct clients to a specific server. This allows clients to return to the same server regardless of the changed TCP/IP connection information. This feature can be very useful when application state management does not leverage a centralized backend store such as a SQL database. In order for content switching to work with SSL, the SSL connection needs to be terminated at the device before the content switch. If this is not done, the content switch fails because the information required is encrypted and unreadable.

Given the complexity and custom nature of many web applications, the IDC design assumes a SQL back-end state management solution and relegates these issues to the specific implementation. The following section describes IDC load balancing in detail.

### 5.5.4 Network Load Balancing

Network Load Balancing is a networking service provided with Windows 2000 Advanced Server, Windows 2000 Datacenter Server, and Microsoft Application Center 2000. The Internet Data Center design requires all services to be both highly available (for example, a redundant server for every service) and highly scalable (both vertically and horizontally). To accomplish this, the design uses two types of clustering:

- Windows Clustering, which may be used to provide fault tolerance for writable services (such as Microsoft SQL Server 2000)
- Network Load Balancing, which provides system fault-tolerance (for non-writable services such as Web servers and servers running Microsoft Internet Security and Acceleration (ISA) Server) and load balancing (for a variety of services). This allows the system designer to horizontally scale network services.

The Internet Data Center design uses Network Load Balancing for three services:

- To horizontally scale and provide fault tolerance for the Web servers
- To horizontally scale and provide fault tolerance for the servers running ISA Server
- To provide fault-tolerance for the management virtual private network (VPN) servers

### 5.5.5 NIC Teaming

In a highly available solution, such as IDC, NIC teaming is typically used. NIC teaming allows two Ethernet network cards to share an IP address. This affords the system designer network fault tolerance due to the ability to connect each network card to a different switch, while providing seamless failover for both the Ethernet port and the switch. Without deploying a NIC teaming solution, IDC would be prone to loose 50 percent or more of the total transactional capacity should a particular switch fail. This maybe an acceptable situation for some customers; however, the IDC design assumes that it is not. NIC teaming is a hardware/driver combination provided by the individual hardware manufacturer, not by the operating system. Features available in the NIC teaming solutions will therefore vary from vendor to vendor. One important option for NIC teaming supported by most vendors is the choice of failover or aggregate operational mode. In failover mode, only one physical interface will be active, whereas in aggregate mode both NICs will be used, thus boosting throughput. In IDC design, failover mode has been selected. The main reason for this choice is predictability. In aggregate mode, when a switch failure occurs, the total network bandwidth for each server is reduced by half. Although the IDC would benefit from a boost in performance during normal operation, after a failure the IDC would run in a degraded mode. Aggregate mode thus compromises the predictability of total system performance, whereas failover mode does not. All of the testing and scalability models in the IDC are performed using failover mode so that its architectural characteristics remain unchanged. In some situations, a switch failure can lead to a cascade effect whereby other elements become overloaded as well. This further reduces the predictability of the total architecture. Any deviations from the prescribed IDC approach should be weighed against the potential drawbacks described above.

Network Load Balancing provides both fault tolerance and horizontal scalability for servers. Network Load Balancing may or may not work properly with NIC teaming, so the combination of the two should be tested thoroughly. Testing is necessary because Network Load Balancing derives its own media access control (MAC) address and assigns it to all Ethernet cards in the cluster. These nodes share the MAC address. NIC teaming typically works the same way, where it derives its own shared MAC address for the network card pair. In some circumstances there may be a conflict between the virtual MAC addresses used by NIC teaming and Network Load Balancing when used together.

The Internet Data Center design successfully uses NIC teaming with Network Load Balancing in unicast mode. However, because the NIC drivers tend to control the team's MAC address and do not detect the MAC address that Network Load Balancing has configured for the cluster interface, a manual change was required. In the NIC configuration tool, the MAC address of the NIC team had to be manually set to the MAC address that Network Load Balancing assigned to the cluster interface.

This configuration has certain limitations. If a new node is ever added to the Network Load Balancing cluster, the NIC teaming MAC address on the new node will need to be set manually. If the IP address of the cluster ever changes, it will cause Network Load Balancing to derive a new MAC address for the cluster, which must then also be replicated across all nodes.

Another important note is how these behaviors interact with Application Center, if it is implemented as part of the content management solution. Typically Network Load Balancing settings are configured on the Application Center cluster controller only. The settings are then replicated out to the remaining nodes of the cluster. If you have configured NIC teaming and then added the node to the Application Center cluster (which will cause the Network Load Balancing interface's MAC address to change to that of the cluster controller), the Network Load Balancing MAC and the NIC teaming MAC will be different. This will cause the addition of the node to fail. To address this, in the Internet Data Center design, Network Load Balancing is configured manually on each node to ensure that the NIC teaming MAC address mirrors the Network Load Balancing MAC address before the node is added to the Application Center cluster.

NIC teaming affords the design the ability to remain at 100% availability in the event of a switch failure, because each network card in the team is connected to a different switch. An alternative to NIC teaming hardware in this situation would be to split Web servers between two switches on the same spanned VLAN. Although this solution would provide redundancy, the design would lose half of the nodes in each cluster in the event of a switch failure. For this reason alone, the NIC teaming solution is superior.

## 5.6 Distributed Load-Balancing Behavior

Unlike hardware load balancers that direct traffic to the nodes that make up the cluster, all (software-based) Network Load Balancing nodes receive every packet destined for the cluster. Essentially, a switch acts like a hub in that it will send traffic to all ports in a particular VLAN on the switch. In unicast mode, Network Load Balancing sends out a decoy MAC address that doesn't really exist. This way the switch never detects which port the MAC address is associated with and therefore sends every packet out to all ports that do not have a MAC associated with them (for example, all ports that have servers that are members of Network Load Balancing clusters).

The switch detects the port associated with MAC addresses for servers that are not members of a Network Load Balancing cluster, so these ports do not see Network Load Balancing traffic. If an additional Network Load Balancing cluster is attached to the switch on the same VLAN as another cluster, the switch sends packets to the ports of one cluster with traffic destined for the other cluster because the switch is not able to detect the ports for any of the nodes of any of the clusters. Without any mitigation, this would be the case in the Internet Data Center network, where there are two separate clusters attached to the same switch with half of the Web servers associated with one cluster and half associated with the other.

To optimize network utilization, packets destined for a cluster should be confined to the ports to which that cluster's servers are attached.

There are two ways to achieve this:

- **Attach all nodes in the cluster to a hub and then attach the hub to the switch port**. In this configuration, a registry change must also be made on all nodes so that a valid MAC address is sent out. This way the switch detects the port and the flooding occurs at the hub.
- **Configure a separate VLAN for every Network Load Balancing cluster**. This way every node in the cluster detects the traffic, but nodes outside of the cluster do not. Although this adds some complexity to the designfrom the VLAN management standpoint, it does eliminate the need for a hub. This way the nodes can take advantage of 100 Mbps full-duplex operation.

The Internet Data Center design uses the second of the two options because it affords maximum network throughput per cluster node (100 Mbps with full-duplex) instead of the reduced throughput which would occur with the presence of a hub, where all nodes would share a maximum throughput of 100 Mbps. This allows the total transactional capacity of the IDC at the network layer to be maximized.

**Note:** Unlike Windows Clustering, which uses a separate interface for the heartbeat, the Network Load Balancing heartbeat is sent and received on a clustered interface. This means that any traffic related to the cluster (inbound client traffic and heartbeats) takes place on the same interface. This configuration cannot be modified. Although some of the Internet Data Center servers that use Network Load Balancing are multi-homed, this is purely for security and traffic segmentation purposes.

### 5.6.1 Setting Network Speed and Duplex Operation

Due to the myriad of combinations of switches and NICs, the port speed/duplex auto-detect feature of most network hardware can prove to be unreliable. To avoid any reliability issues associated with this, all NICs and all switch ports are forced to 100 Mbps full-duplex operation by using the manual settings of either the hardware or software drivers.

### 5.6.2 Network Load Balancing in Unicast Mode

Network Load Balancing operates in one of two modes: *unicast* or *multicast*. Using multicast mode requires static Address Resolution Protocol (ARP) entries at the router, and can cause additional processing at the switch, therefore all Network Load Balancing cluster nodes in the Internet Data Center designare configured for unicast mode. However, when using Network Load Balancing in unicast mode, all nodes in a cluster are unable to ping other nodes in the same cluster on the clustered interface, although nodes outside of the cluster are still able to ping each node. This occurs because the members of a cluster's clustered interfaces share a MAC address. When a node uses ARP to find another node in the cluster, TCP/IP returns the originating node's MAC address (which is shared amongst all nodes). This results in the request never leaving the node, because ARP causes the originating node itself to be resolved to the MAC address. If the nodes within a cluster are multi-homed they will still be able to ping each other through non-Network Load Balancing interfaces, however.

### 5.6.3 Using Network Load Balancing on Web Servers

Although Network Load Balancing can technically support 32 nodes per cluster, in practice, it is best to keep cluster sizes at approximately twelve or less for the following reasons:

- Application Center 2000 allows a maximum of twelve nodes in a cluster.
- CPU overhead to process the load balancing algorithm grows with each additional node that is added to the cluster.
- More than twelve nodes may become unmanageable for routine content deployment or system maintenance.

In the Internet Data Center design, two Network Load Balancing web server clusters are configured that could each be scaled to twelve nodes within each cluster. Two private virtual IP addresses are assigned to each internal Web cluster (one for each) in which the perimeter ISA firewalls will redirect external web traffic to these internal web clusters. Each time a request comes in from the Internet, ISA will apply a random number algorithm to determine which internal web cluster to forward the request to.

Other firewalls may require a round-robin DNS implementation to maintain high availability across each Network Load Balance web cluster. This configuration would maximize the benefits of both technologies and remove the shortcomings of just implementing round-robin DNS. Two public IP addresses would be assigned to the same host name, such as www.northwindtraders.com, in which the perimeter firewall would redirect external requests from the Internet mapping each public IP address to each internal web cluster private address in a round-robin fashion.

Designs which implement either ISA firewalls or third party firewalls using round-robin DNS both maintain high availability. Since each network card in a NIC team configuration is connected to a separate switch, a switch failure would not impact the site since both web clusters would still be active. Both switches would have to fail to make the web clusters, and therefore the site, completely unavailable.

Using round-robin DNS as the sole load-balancing solution (without Network Load Balancing) has the following major deficiencies:

- If a node in a cluster fails, the node's IP address would need to be manually removed from the DNS table. There is no dynamic removal.
- DNS servers on the Internet may have the entry of the failed server cached, and would give that address to clients until the Time to Live (TTL) on that entry expired. This could be hours or days.

Due to these deficiencies, round-robin DNS does not make for a very good load-balancing solution on its own. When round-robin DNS is used in conjunction with Network Load Balancing, however, those deficiencies no longer exist.

### 5.6.4 Network Load Balancing and Ports

The following Network Load Balancing port rules are configured on the Web servers:

| **Rule** | **Start** | **End** | **Protocol** | **Mode** | **Load** | **Affinity** |
|---|---|---|---|---|---|---|
| 1 | 0 | 79 | Both | Disabled | N/A | N/A |
| 2 | 80 | 80 | TCP | Multiple | Equal | None |
| 3 | 81 | 65535 | Both | Disabled | N/A | N/A |

*Source: Microsoft, 2002* Table 2: Web server Network Load Balancing port rules

**Note:** The default port rule of single affinity for ports 0-65535 was removed.

Unlike the ISA Server configuration, the Web servers only need to process requests on port 80 (HTTP). Normally port 443 would be configured for single affinity to the Web servers. It is not, however, because the HTTPS connection from clients is terminated at the ISA Server computers that make up the Internet perimeter firewalls. The ISA Server computers request the information via HTTP from the Web servers and then send it back to the client over HTTPS. This is terminated at this firewall for two reasons:

- It gives the servers running ISA Server the ability to inspect the contents of the transactions, rather than just passing them through encrypted.
- It allows the use of "no affinity" with Network Load Balancing rather than "single affinity."

If HTTPS were allowed to pass through the ISA Server firewall to the Web servers, single affinity would need to be configured for port 443 on the Web servers. This would cause Network Load Balancing service to load-balance on the source IP address only, rather than on both source IP address and port, as it does when affinity is set to "None". ISA Server Web Publishing sends the ISA Server computers' internal interface IP address as the source IP address, not the clients' source IP address. In single affinity mode, if the clients' source IP addresses were being sent, there would still be good load-balancing as the number of source IP addresses would be large. However, when using Web Publishing, the number of possible source IP addresses is equal to the number of ISA Server computers. This would cause HTTPS traffic load-balancing to greatly diminish, because most requests would end up only being serviced by a few Web servers.

Since HTTPS traffic will be terminated at the perimeter firewall, "Force HTTPS" should not be turned on in the Web site or server level properties. To make the client use HTTPS, the Web pages should be coded to reference a page as https://, rather than simply linking to a page. Additionally, when the secure portion of the transaction is complete, hard-coded http:// links should be provided that force the client browser back into standard HTTP mode. Although Web application designers do not typically like to hardcode links, this should be done to move the user in and out of HTTPS mode.

If the Web application were using HTTP exclusively but also maintaining state at the Web server by using, for example,  ASP session state, HTTP would also need to be set for single affinity producing the same effects as using single affinity on the Web servers for HTTPS. To address this, the Request Forwarder functionality of Application Center can be used. However this feature does incur additional resources on the Web servers so maintaining session state on the Web server is not recommended or used in the Internet Data Center design.

### 5.6.4.1 TCP/IP Settings

Both a dedicated and cluster IP address are configured on the Internet facing interface. This allows monitoring tools to be used to test response time for each node in the cluster directly. This way, if a node becomes unresponsive, the monitoring solution could remove it from service. Without a dedicated IP address on these interfaces, only the response time of the cluster could be measured, and this would not provide the information necessary.

### 5.6.4.2 Using Network Load Balancing on VPN Servers

The VPN servers in the base Internet Data Center design use Network Load Balancing as the high availability mechanism. This is because in this design the VPN servers are used for remote management which means they are unlikely to handle more than 20 simultaneous connections at any given time. Each of the ISA Servers used is perfectly capable of handling this traffic; therefore, the load balancing and scaling out features of ISA Server are a low priority. However if the design will be extended to use these servers as gateways for your application logic, say as part of a business to business (B2B) solution, load balancing and scaling out can be very important.

**Note:** The Point-to-Point Tunneling Protocol (PPTP) requires the setting of single affinity due to the session orientation of VPN connections.

### 5.6.5 Port Rules

The use of ISA Server as the VPN solution in the Internet Data Center design means that ISA Server handles the additional port filtering security. In this case, the default port rule was left in place. Although it is not required to configure all ports, they are left in place for the sake of simplicity. The ISA Server still only listens on the appropriate VPN ports, so the additional complexity of managing granular port rules in two locations is avoided:

| **Parameter** | **Setting** |
|---|---|
| Port Range | 1-65535 |
| Affinity | Single |
| Filter Mode | Multiple hosts |
| Load | Equal |
| Protocols | Both |

*Source: Microsoft, 2002* Table 3: ISA Server VPN Network Load Balancing configuration

Single affinity is used because the users of the VPN solution are typically working from home, where they are likely to come through proxy servers at their ISPs.

Class C affinity would serve no purpose because it addresses the problem of clients going out through multiple proxy servers at an ISP. As VPN connections require a session, if a user goes out through multiple proxy servers at their ISP, the session would break at the ISP's proxy server before getting to Network Load Balancing on the VPN servers.

#### 5.6.5.1 TCP/IP Settings

The Internet Data Center VPN design does not use a dedicated IP address on the clustered interfaces. It uses only a cluster IP. The reason for this is primarily security related, but is also for simplicity. With no dedicated IP address configured, these servers are not addressable individually by the Internet. They can only be addressed as a cluster and therefore a client never knows which server answered its request. So if a hacker were able to somehow breach one of the servers and place a piece of Trojan software on it, for example, they would not know which node it was placed on and therefore would be unable to address the Trojan directly to launch it. While that certainly does not rule out the possibility that the hacker might get load-balanced back to that node, it does mean that it would involve some additional persistence.

## Conclusions

This option also has an advantage of being accessible via the Internet to support staff wherever they are. As long as they can provide the correct certificate to authenticate with the VPN solution, support staff will be able to obtain a secure connection to the Internet Data Center network.

Internet Data Center design is capable of supporting both remote facility options as well as the local dedicated VLAN solution.

# CHAPTER 6

# CONCLUSIONS

The design of Internet Data Center (IDC) conceptual model provides a model to be used for implementing a comprehensive solutions in a company generally. The design is evaluated based on the Network Scalability, Network Availability, and Security Manageability.

The current Internet Data Center design completely locks down all Web servers by using a Web server security policy and Active Directory organizational units.

Since the Web servers in the Internet Data Center design are multi-homed (two network interface cards, or NICs), architects were concerned about hackers gaining access to the production network through the back-end network interface card (NIC). The design adds another layer of protection by separating the DMZ VLAN from the rest of the production VLANs by placing a firewall directly between the internal network interface of all the servers in the DMZ VLAN and the other internal VLANs. All traffic from the DMZ VLAN that flows to the servers that are in the production VLANs must go through the firewall first. If hackers did gain access to a Web server, they would still need to beat the security configuration of the internal firewall before they could damage data.

Having a separate data and management VLAN allows the most important servers (normally the servers running SQL Server) to be placed behind two sets of protection. First, the Internet Data Center design uses stateful inspection and firewall access control lists (ACLs) to control the communication of TCP and UDP ports between servers in the DMZ VLAN and servers in the data and management VLAN. Second, the design uses VLAN technologies and additional access control lists on the switch that can be configured to control the communication of TCP and UDP ports between the infrastructure VLAN and the data and management VLAN.

Network availability can be achieved by providing redundancy at every level and by using automatic failover. Two network devices are implemented within each layer of the design to provide high availability at the network level. Duplicate routers, switches, and firewalls are implemented to maintain availability throughout the network. There is no single device in the design that would bring the site down. If the firewall fails, a backup firewall takes over. If one switch fails, another one takes on full load until the first one is repaired. If a Web server's network adapter fails, another NIC becomes active automatically with no impact on traffic flow. If a complete Web server fails it can be taken offline, repaired, and added back into Web farm without any impact on production. The database partitions on the SQL Server computers are protected as part of a SQL Server database cluster.

Network traffic is becoming more and more unpredictable. The old 80/20 rule held that 80 percent of network traffic was limited to the workgroup, with only 20 percent involving the Internet. But with the increasing use of e-business systems, the current ratio is closer to 50/50. If trends continue the ratio may invert to 20/80, significantly increasing backbone traffic. As the Internet backbone bandwidth increases, it will increase network demand on the e-commerce sites.

Technological development is moving fast to provide technology that will ease the pressures in e-commerce networks and provide a path for upgrade to higher bandwidth requirements. The network design should include new technologies, such as Layer 2 and Layer 3 devices that switch and route traffic at wire speed. Modular and stackable switches offer port density and port speeds up to 100 megabits per second (Mbps). These devices also provide solutions for e-commerce data centers where the switch can be stacked with gigabit Ethernet (1000-Mbps) links, and provide thousands of high-speed ports.

Bandwidth aggregation for servers is available through multiple adapter technologies, which eliminate server bottlenecks by allowing incremental increases in bandwidth between a server and a switch. These technologies enable high-speed transmissions that extend the capacity of the physical medium.

## REFERENCES

Goodyear, M. Editor (1999). *Enterprise System Architectures*: CRC Press.

Rajput, W. (2000). *E-Commerce Systems Architecture and Applications*. Artech House.

Chen, Q., Hsu, M., Dayal, U (2001). *Peer-to-Peer Collaborative Internet Business Servers*. HP Laboratories Technical Report HPL-2001-14.

Kan, G. (Editor) (2001). *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. O'Reilly & Associates

Peraire, C., Coleman, D. (2000) *Modeling for E-Service Creation*. Technical Report.System Design Laboratory, SRI International

Kotov, V. (2000) Towards *Service-Based System Organization.* . HP Laboratories Technical Report to be published.

Snevely, R (2002). *Enterprise Data Center Design and Methodology*: Sun Microsystem Press.

Dodds, T (2000). *Network Security Fundamentals and Overview* : Prentice Hall Inc.

McReynolds, M (2001). *Planning and Building a Data Center*, Sun Microsystem Press.

Cassidy, J.C (1998). *Data Center Architectural:* CRC Press.

Louis, G (2002). *Strategies for Success in the Network Economy*: Sun Microsystem, Inc.

Gergg, M (2001). *Data Center Evolution Strategies*: Prentice Hall.

Cowley, R (2002), *Solving the Problems of Data Center Web Serving*: Zeus Technology

Gerard, B (2002), *Global Infrastructure Risk Management: Planning & Recovering in uncertain times:* Aperture Inc.

Microsoft, (2002), Architectural Elements: Available at: http://www.microsoft.com/idc

ISA Server firewalls (2002). Available at: http://www.microsoft.com/isaservers

Integrated Value-added Services for Internet Data Centers. Available at: http://www.cosine.com/idc

Choose the Best Data Center Co-location Outsourcing Solution for your e-Business. Available at: http://www.certsolutions.com

Emigratus Data Center Featured in Entrepreneur Magazine. Available at: http://www.emigratus.com/idcent

Facilities Security Key for Data Center Selection. Available at: http://www.webhostingindustry.com/data_center_selection.html

Instant Data Center. Available at: http://www.instantdc.com/web_hosting.asp

Internet Data Center Web Hosting. Available at:
http://www.microsoft.com/business/services/mcsmsf.html

Internet Data Center Management and Operation System. Available at:
http://www.microsoft.com/business/mos/srfqty.html

High Performance Data Center. Available at:
http://www.upsite.com/exchange/datacenter.asp

Planning and Building a Data Center. Available at:
http://www.intel.com/idc

International American Engineering Association (IAEA),1993.Available at:
http://www.deakin.edu/conpmodel/iaea

The Forrester Report: "How to buy web infrastructure"(June, 2001): Available at:
http://www.idc.com