Planning and Building a Data Center

Meeting the e-Business Challenge

Executive Summary

As the Internet access market matures, Internet Service Providers (ISPs) are facing significant competition and challenges to profit. The businesses that will succeed in this market must either:

- Achieve the economies of scale necessary to support a low price business model; or
- Offer added value, typically in the form of specialized services such as applications hosting to justify a premium price.

Both of these goals require the right facility and management.

Many service providers are attempting to make this transition but lack sufficient information on how to develop and implement a data center of such scope. This document provides a high-level overview of the requirements for successfully establishing and operating an Internet data center in today's marketplace. It offers some of the key steps that need to be taken, including project definition, prerequisites and planning.

Table of Contents

Executive Summary . . . 2
Meeting the Market Challenge . . . 2
Assumptions and Limitations . . . 3
Service Definitions . . . 3
Managed Services . . . 3
Co-location Services . . . 4
Advanced Services . . . 4
Example: Application Service Provision . . . 4
Infrastructure: Layout . . . 5
Building Layout . . . 5
Operation . . . 5
Internal Layout . . . 5
The Technical Suite Concept . . . 5
Secure Vaults . . . 6
Racking Neighborhoods . . . 6
Private Cages . . . 6
Center Capacity . . . 7
Facility . . . 7
Power Specification . . . 7
Resilience . . . 7
Building Security and Access Control . . . 7
Fire Control . . . 8
Air Conditioning . . . 8
Staff Facilities . . . 8
Systems Infrastructure . . . 8
Rack Configurations . . . 9
Data Center Management and Operation . . . 10
Service Management Center (SMC) . . . 10
Service Monitoring and Maintenance . . . 10
Customer System Backups . . . 11
Storage Area Networks . . . 11
SAN and Network Attached Storage . . . 11
Problem Management, Configuration Management and Change Control . . . 11
Data Center Organization . . . 12
Structure . . . 12
The Production Environment . . . 12
The Development Environment . . . 14
Conclusion . . . 14
General . . . 14
Product Balance . . . 14
Integration . . . 15
For More Information . . . 15

Meeting the Market Challenge

Today, many service providers are aiming to differentiate themselves and increase profitability by offering a range of Web-based services targeting business customers. At the same time, the market for outsourced Web services is growing rapidly as small- and medium-sized businesses seek to migrate from traditional hosting options—such as co-location—to newer, service-focused hosting methods. Forrester Research has projected that this market will grow from $152 million at the end of 1998 to more than $5.3 billion by 2003.

Together, these factors have led to the emergence of a new type of business provider in the Internet marketplace. These service providers have moved beyond traditional Web hosting and access to more diversified, complex Web-based business services. One such example is the Application Service Provider or ASP. ASPs are companies that provide services to deploy, host, manage and rent access to an application from a centrally-managed facility.

The ability to provide more robust services to customers raises significant new challenges for traditional service providers in the areas of performance, scalability, availability and security. Adding to these challenges is the unpredictability of such a dynamic market. The service provider must be able to manage, transparently deliver and back-up vast quantities of data for a wide range of clients. The infrastructure must be able to scale quickly, and when it comes to providing more comprehensive services (such as applications), a significant additional investment in support and implementation skills is required.

In order to construct a data center that can meet the challenges of the new market, there are three basic areas of data center definition and development:

- **Facilities:** including building, security, power, air-conditioning and room for growth
- **Internet connectivity:** performance, availability and scalability
- **Value-added services and the resources to support their delivery:** service levels, technical skills and business processes

The aim is to provide customers with the physical environment, server hardware, network connectivity and technical skills necessary to keep Internet business up and running 24 hours a day, seven days a week. The ability to scale is essential, allowing businesses to upgrade easily by adding bandwidth or server capacity on demand.

Item	Standard	Optional (chargeable)
Network Monitoring	Y	-
Server Monitoring	Y	-
Web Server	Y	-
E-mail Service	Y	-
Hardware Support[1]	Y	-
Backup[2]	Y	-
Remote Management[3]	Y	-
Web Usage Analysis/Monitoring[4]	-	Y
Firewall	-	Y
VPNs	-	Y
ISS	-	Y
Load Balancing[5]	-	Y

For custom solutions, it is recommended that hardware support should only be provided if the following conditions are met:

1. The data center supplies and manages the hardware (i.e. the customer cannot supply additional hardware).

2. A backup system is in place.

3. Remote management is provided via a recognized software application and is limited to the facilities provided by that application.

4. A standard analysis tool is implemented.

5. Load balancing is only offered after a suitable hardware solution has been evaluated and tested.

***Table 1:** "Standard" and "Optional" Maintenance and Support Options—Managed Services*

Assumptions and Limitations

It is assumed that the reader is already familiar with the basic business planning and investment requirements that must be met when setting up a data center. Therefore, while the commercial objectives in establishing a data center are considered by this document, it does not probe the specific areas of business planning or investment. Rather, this document assumes the reader is seeking input on the three major areas of data center definition and development, namely:

- **Facilities:** building, security, power, air-conditioning, growth, etc.
- **Internet connectivity:** performance, availability, scalability etc.
- **Value-added services and the resources to support their delivery:** service levels, technical skills, business processes, etc.

For the purposes of this document, the report is written for a data center which has a total floor capacity of between 1,000 and 30,000 square feet of space and is operated by a team that specializes in facilities management.

Service Definitions

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

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.

Table 1 contains a matrix of the typical maintenance and support elements that can be supplied as a part of standard managed service, together with a list of possible optional items.

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.

Advanced Services

Example: Application Service Provision

Front office applications are ideal candidates for the ASP deployment model. Most companies implementing front office systems need to serve a geographically 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 infrastructure 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.

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.

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.

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.

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.

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 below.

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. 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

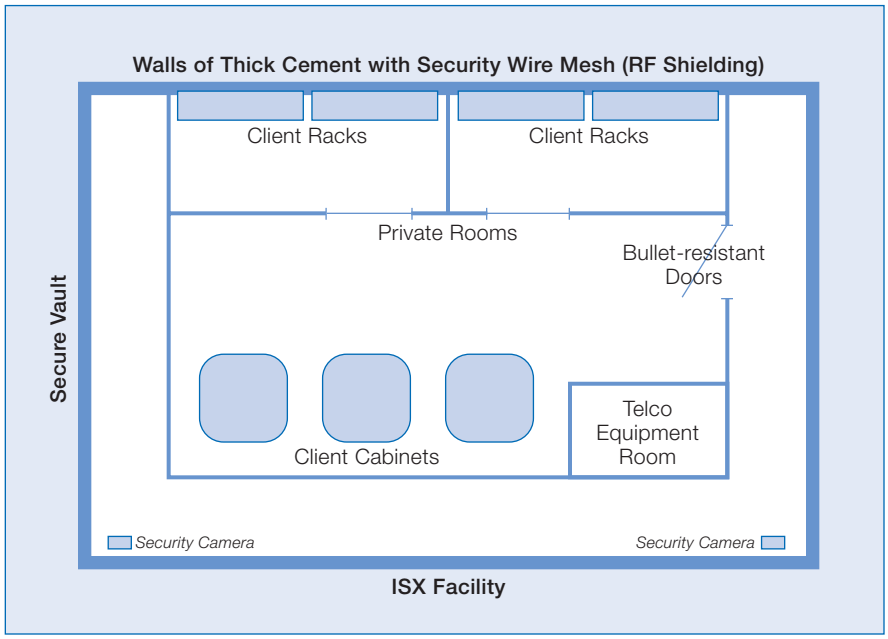

Figure 1: *Typical Technical Suite "Secure Vault"*

- 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.

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. A typical secure vault is shown in Figure 1.

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. 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

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.

Technical Suite	Servers / 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 2: *Server and Rack Capacity per Floor.*

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.

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.

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, UPS-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 Table 2.

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.

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.

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.

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.

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 68º F, plus or minus 3º F, with humidity at a constant 45%-50%. Cooling units should be placed over walkways or hallways, never over racks.

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.

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 bandwidth 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.

Designed to accelerate and improve the overall performance of the data center, Intel offers the following products:

Intel® NetStructure™ 1500 Cache Appliance
This Internet Caching Appliance is designed to enable caching solutions that can significantly alleviate bandwidth constraints, improve access and boost Internet performance.

Intel® NetStructure™ 7110 e-Commerce Accelerator
This drop-in, scalable appliance processes up to 200 secure cps, improving Web site performance by up to 50 times by offloading encryption processing.

Intel® NetStructure™ 7170 Traffic Director
This advanced Layer 4-7 traffic management and load balancing solution enables priority error-free service levels to be set for specific customers or URLs for improved site performance.

Intel® NetStructure™ 7190 Multi-site Director
This comprehensive traffic management solution offers a choice of efficient traffic routing methods so that requests may be sent to the fastest-responding site or to the best site for optimal performance.

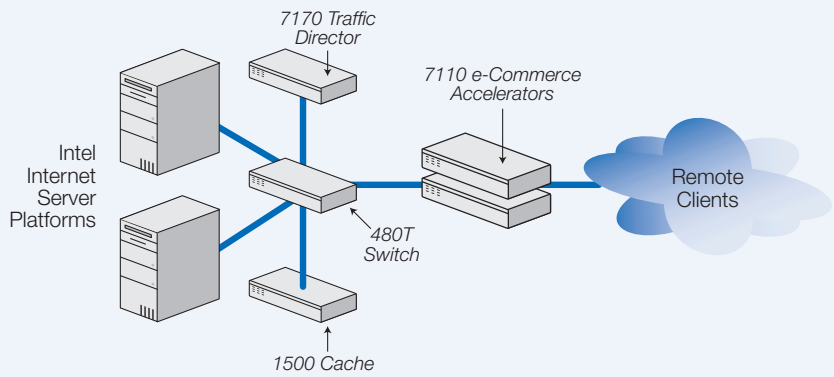

Server Type	Number of Switch Racks per Neighborhood	Number of Server Racks	Number of Network Connections per Customer/Server
Managed Services	2	14	3[1]
Co-Location	0[2]	18	1

Notes: Of which two of these connections are for management network connection.
Because only one network connection is provided, the number of switch racks in a co-location technical suite is typically very small. These are commissioned separately from the server rack neighborhood.

Table 3: *Neighborhood Racking/Network Configurations*

Intel offers a variety of products designed to provide high reliability and flexibility in the data center:

Intel® ISP1100 Internet Server

Engineered for quick, high-volume deployment, the ISP1100 rack server features a high-density 1U design. Support is provided for either the Intel® Pentium® III or Intel® Celeron™ processor, so performance can be scaled from less demanding to more robust applications. This is an ideal platform for dedicated hosting.

Intel® ISP2150 Internet Server

With support for up to two 700 MHz Intel® Pentium® III processors, the 2U ISP2150 provides the performance needed for demanding Internet applications such as secure Web and CGI scripts. The ISP2150 can be clustered, load balanced and managed from a central console.

Intel® NetStructure™ 6000 Switch

A high-density gigabit and fast ethernet data center switch, this product is designed for maximum reliability when aggregating servers. It includes redundant fans, power supplies and control processors.

Intel® NetStructure™ 480T Routing Switch

A medium density gigabit and fast ethernet data center switch. Traffic shaping features enable the service provider to control the bandwidth going to and from servers. Customers can pay depending on the amount of bandwidth they require.

Intel® Express 550T Routing Switch

Low-density fast ethernet data center switch. Transparent cache switching feature allows the service provider to cost-effectively integrate Cache Appliances at a fraction of the cost of a traditional layer 4 switch.

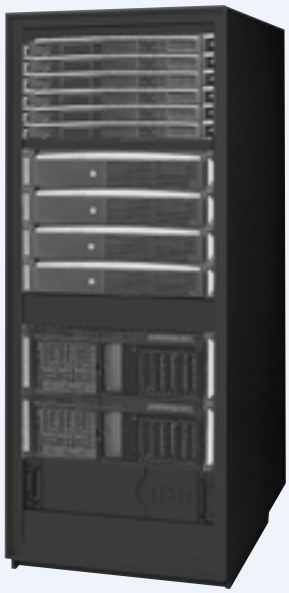

Within the data center, the network infrastructure should be based on Gigabit topology, which offers scalability, resilience and manageability. Another proven technology, ethernet is economical, highly scalable (10/100/1000 MB) and has a large user base. This means there are many vendor solutions available,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

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. As referenced in Table 2 (page 6), 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. Here, Table 3 provides

a general guideline for the ratio of server racks/ switch racks and network connections per neighborhood.

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. Since the Intel® Architecture offers a broadly compatible, open server platform, it is an ideal choice. Unlike proprietary designs, the flexible architecture of Intel® Internet servers provides a wide selection of choices at every level of the solution stack.

Data Center Management and Operation

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.

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)

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.

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[1]. 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, automated 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 (i.e., all tasks are charged at a minimum of $200 per hour).

[1]In order to ensure high quality managed services, the ability to monitor and control data center operations should extend to the server hardware itself.

[2]OpenFile was originally developed as an activity monitor utility to provide a quick and convenient means of identifying all open files on a VS system. The application has been greatly extended to run on most operating system platforms to provide a means of ensuring open files are safely backed-up.

Customer System Backups

Backups should only be provided for Managed Services. 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*[2] be evaluated as a part of the backup solution.

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.

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

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, as illustrated by the following example:

A customer reports an incident—a server is not responding. The call is logged via the CMS and attempts are made to resolve the incident, but changes are required to the server. The incident then becomes a problem that may impact other servers as well. The problem is rectified, and as a result, a change is logged in the Change Management System recording the fact that the configuration of the affected server has been updated. The change is used to update the relevant server details in the Configuration Management Database. At the same time, the CMDB is interrogated to find out if other servers may be affected by the problem and need their configurations changed. These servers can then be updated.

Because of the way they are interrelated, 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)

Data Center Organization

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

As shown in **Figure 2** (page 13), the 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 feedback, 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.

The Production Environment

The production environment is subdivided into three areas: pre-sales, project management and the Service Management Center.

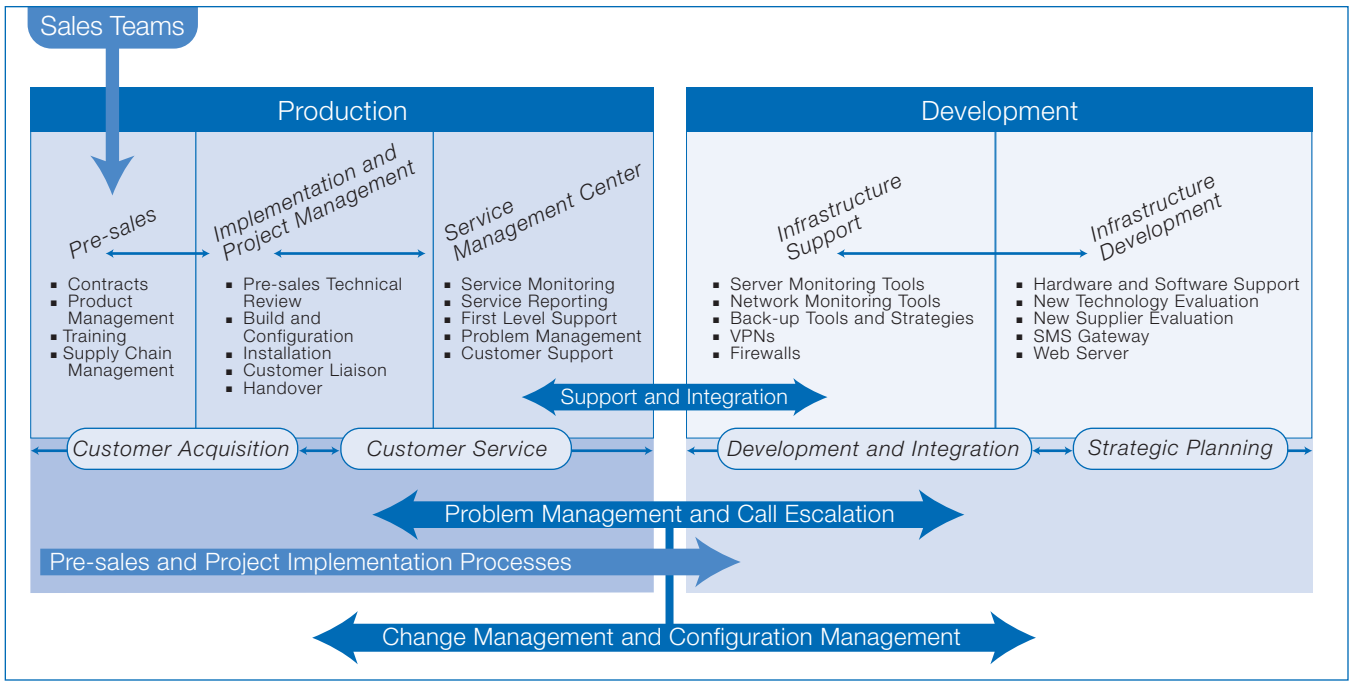

Figure 2: Internet Data Center Organization

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.

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.

The Development Environment

The development environment is subdivided 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.

General

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.

Product Balance

Careful consideration must also be given to the balance of products—Managed Services against Co-locate services; customized managed services against "off-the-shelf" Managed Services, where all the support elements, etc., are pre-defined.

Co-locate services offer a quick means to earn revenue—the racking space, power and network link can be charged at a premium, but there is no overhead incurred in maintaining the services.

Managed Services offer the most constructive way of generating income, especially if a range of e-Commerce applications can be supplied to clients as a part of the baseline service, providing them with a value-added resource.

Integration

As defined in this white paper, the real secret to success within the Internet data center/ASP environment is that of integration. It is not sufficient to have a robust set of products (servers and software) and an environment in which to operate them. A successful data center implementation is one that presents customers with a seamless, fully-integrated environment—from the initial point of sale, through to on-going management and support.

The data center market is still relatively young, and demand is growing. Intel has a product and services suite that is ideally suited to the ISP/ASP hosting environment, offering as it does a flexibility of approach and implementation that can be used to structure a market-leading set of services for potential customers and clients. By combining these types of products with operational recommendations as given throughout this white paper, service providers can successfully meet the challenges presented by today's rapidly changing and competitive marketplace.

For More Information

For more information about the emerging e-Business data center, visit: http://www.intel.com/ebusiness

Intel offers a comprehensive range of products designed to provide maximum flexibility for data center environments. To learn more, see: http://www.intel.com/network/products

Intel ISP Program provides dedicated resources and Internet-focused solutions that can help you grow your business. Sign up today at http://www.intel.com/isp.

NP1654