

The Internet Delivery Stack

A Framework for Successful Internet Computing and Applications

January 2008

Audience

Business executives, founders, and chief technology experts of companies facing new Internet based initiatives or recovering from a failed infrastructure.

Abstract

Internet applications are actually complex software applications that require many pieces to work properly to deliver a successful experience to the end-user. This paper provides a framework – called the Internet Delivery Stack – to plan and track the myriad of components necessary to make it all work, and presents how Online Tech delivers its co-location and dedicated server solutions against this framework. Using this framework you can plan and measure the security, scalability and reliability of your Internet delivered applications and the IT infrastructure supporting it.

Introduction

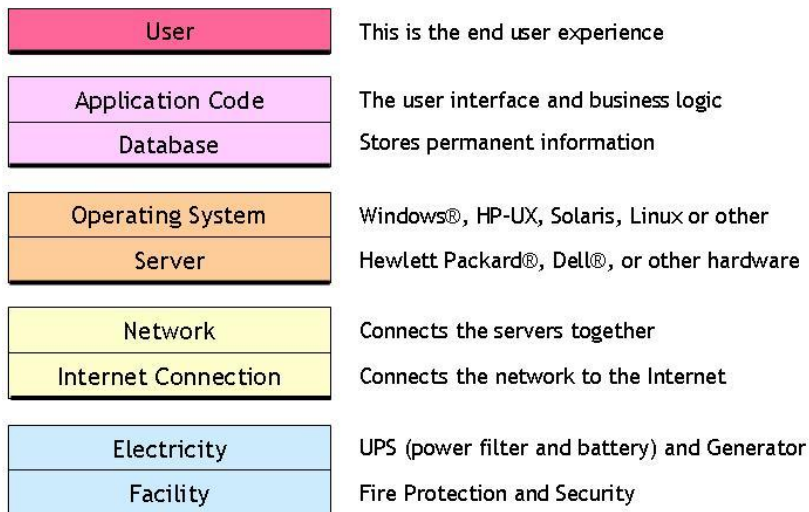
The Internet is assuming an ever increasing critical position in business. Whether it is used to market your company; take or process sales; track customers; deliver e-mail, audio, video, data or other media; serve customers; or deliver your entire product—the Internet is assuming a more important role each day—certainly not a less important role each day.

An Internet-based service is expected to be *secure* (can't be hacked, data is secure), *reliable* (doesn't go down, if it does there are fail-over scenarios) and *scalable* (if we add 10,000 users we don't have to re-architect everything). So, what does it take to assure a success on the Internet? It takes many complex systems working together. The "Internet" is a complex web of software, hardware, service and infrastructure that must all work together in a highly fluid environment – who knows exactly what the 1,000s of potential users, will actually do?

Different applications have different standards to meet for each of these facets relative to the system components. To maximize each of the three facets (*secure, reliable, scalable*) at all levels is prohibitively expensive even for the largest of ventures. That is why Online Tech has built a framework describing each layer of components enabling you to use the framework as a road map to develop High Availability applications and a Disaster Prevention and Recovery Plan (DPRP) for your critical data and applications so that management can make intelligent trade offs given their specific application and business state.

I. The Internet Delivery Stack

At Online Tech, we have a unique framework we call the **Internet Delivery Stack™ (IDS)** to identify all the elements necessary to deliver a successful experience to your end-users. We use this framework to identify roles and responsibilities and to scope projects. There are nine layers of components to the IDS. Each layer requires designing, implementing and tending to assure a certain degree of security, reliability and scalability in that layer. These layers are described in the diagram below.



Internet Delivery Stack™

A. SECURE, RELIABLE AND SCALABLE STACK

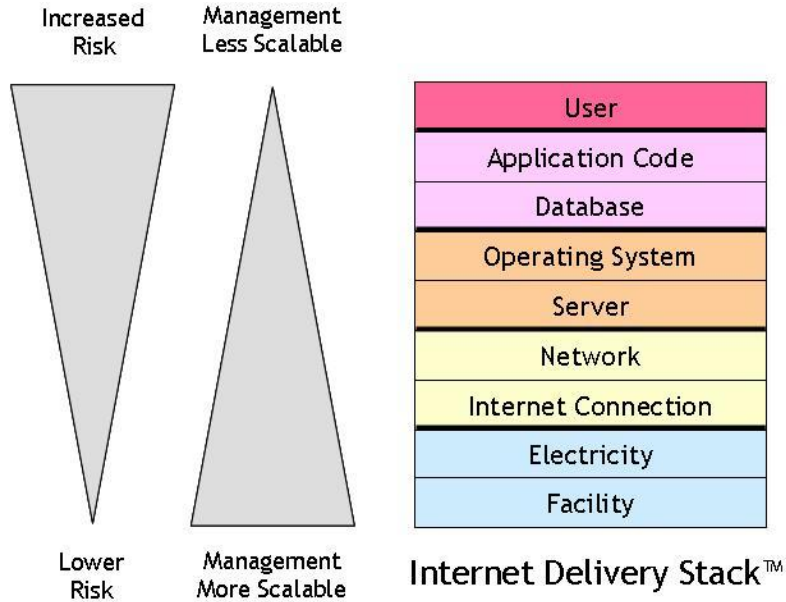
The security, reliability and scalability of the end-user experience will be limited by the security, reliability and scalability of the weakest layer in the stack. So, it's important to understand each layer of the stack. While there are many technologies and services that can increase security, reliability and scalability ultimately much will be determined by the architecture of the entire solution.

It's the sum of the security; reliability and scalability for each layer of the stack that ultimately determine your ability to deliver an experience of service to your end-user.

B. RISK

There is an interesting risk characteristic of the stack. As you go up the stack the layers contain higher and higher degrees of risk and lower degrees of commoditization. For example, there are 100s of people in most 10 square mile areas who can help with electrical problems. There are fewer, but still plenty that can help with Internet connectivity and most simple networks. Even at the hardware and OS layers there are generally dozens of people who can help when things go sour – which they seldom do. In fact, in our experience we can keep a Microsoft® Windows 2003 or Linux® server running almost indefinitely with 100% uptime if there is no application or database running on it. The real risk is in the application code and database. For generally, there are very few

people who know those layers (often 3 to 5 will be the only real experts on the source code or database design).



C. CONCLUSIONS

Internet initiatives are expensive but can also be hyper-successful requiring you carefully match the spending on each layer of the Internet Delivery Stack to the stage of your business. By using this framework to identify where you are today and where you need to be as you grow you can set goals and track progress.

In the following sections, we describe each layer of the stack.

II. End User Experience

User

The end-user experience is pivotal to successful Internet applications. Said another way, whether delivering business processes, sales, marketing, support, media, communications, or applications via the Internet—success requires satisfied end-users. No matter your business or purpose, without the end user, your online service might as well not exist. It is critical for any business using the Internet to define a framework built around meeting and exceeding the end-users needs.

The end-user can be customers, employees, partners, patients, citizens, students, et cetera—clearly, an extremely diverse group of people. It is the end-user's experience we are attempting to control and is the ultimate measure of success.

In winning your end-users' acceptance (trust)—and protecting your company—your business requires a hardened Service Level Agreement (SLA). We back our SLA with what is referred to as a *disaster prevention and recovery plan (DPRP)*. The DPRP is the roadmap: the declaration of facility hardness; the documented process of facility and network architecture; the monitoring and testing; and defines accountability.

The DPRP must be documented and have a formal review process (*hope is not a strategy!!!*). This component allows management to make intelligent decisions given your specific application, risks they are willing to take and the requirements of the business. Without a documented and continuously updated DPRP plan there is no reliable form (proof) of preparedness and no foundation for supporting the SLA. For the *IDS*, it is highly recommended to complete the DPRP from the infrastructure level, through the server and operating system levels, to and including the database/application level. This provides a complete assessment of the risks and the level of service you are delivering to the end user.

III. Application

Application Code

The application is the user interface and business logic. It interacts with the database that contains the permanent information. But as you see with the IDS, the application and database is the unique core of the online solution, but can not operate without the other layers of the framework.

An online solutions delivery's greatest risk resides with the application and database. Why?

- First, typically only a handful of people in the entire world know the code—the developers.
- Second, did the developers create the code with the entire IDS built in to the code's architecture? If the code isn't designed to meet scalability requirements, the entire application may break as the number of users grows. In this case, the entire application may have to be re-architected to work in a load-balanced multi-server environment as the application scales.
- Third, has code monitoring scripts and testing been defined and created to assure a high quality solution and bug-free deployment? Defining the test requirement is usually best done prior to writing the code.
- Fourth, were required server and operating system settings communicated to the server technicians? It's not atypical for the application developers to make assumptions that surprise the infrastructure provider when the application is ready to launch.
- Last, was the code written to be an enterprise solution with appropriate security built in? Security as an after thought is extremely difficult to patch into the system and should be a requirement developed and reviewed with security experts.

For these reasons, it is important to detail and communicate who is responsible for the design, development, support and maintenance of the application itself, including the code and all related scripts. In addition, involving the system operators provides the application development team with services to help sleuth anomalies and to prepare for application upgrades.

Anomaly Research – Often bugs can best be solved by a joint research effort of the application development team and the systems operators. It is strongly suggested to have system operators working with the application team to help sleuth bugs. This might include reviewing performance logs, setting up test environments or reviewing system configurations. For example, the application team may request a systems engineer and network engineer to monitor utilization during certain functions to help identify bottlenecks or crash points.

Application Upgrades – As the application is upgraded it's best to work closely with system and network engineers to be sure the architecture is properly configured for the applications needs. With regular communication between the application team and the systems, database and network engineers you can be most assured that new releases will be trouble free.



At Online Tech, we work directly with our customers' developers; third party developers/certified development partners; or use our own development staff to meet our customers' needs.

IV. Database

Database

The online database's architecture is a critical factor in the ability for an online solution to scale. Many businesses have not had the expertise or insight to appropriately develop an enterprise-class database that can meet the business' growing online needs. The result—the business either failed at delivering the online service or was forced to step back and rebuild the database and application to handle the growth costing significant investment and lost revenue opportunity.

In order to assist our customers, Online Tech offers two additional services beyond developers and technicians. First to help growing businesses succeed, Online Tech brings senior level IT management expertise via our Virtual CTO (vCTO) managed service. Secondly, Online Tech is able to offer database management services and/or on-going support through our Remote DBA service. These services bring the breadth of expertise and resources required to complement or supplement your business' demanding IT needs.

Operating System
Server

V. Operating System and Server

Your online application and database require at least one or multiple servers with operating systems up and running. Your decision in regards to risk and responsibility that you wish to retain or outsource defines which of the following options your business selects.

A. SERVER SERVICES

Businesses have multiple server options to run their website or application from the low cost sharing of servers; to providing their own server; or gaining the flexible service of having a dedicated server provided and managed for them.

- **Shared server hosting** – the lowest cost solution is sharing your application on a server provided by Online Tech and shared with multiple customers have their application or website on the same server. Shared servers are cost effective due to sharing the cost and managed services with other businesses, but it also exposes your website or application to the stability issue of other applications on that server.
- **Hosting your server (Co-location)** – On the other end of the spectrum, you can configure and provide your own server hardware, operating system and utilities and run that server in Online Tech’s data center. Online Tech provides the lower layers of the IDS and you manage the server operation either through direct physical access or remotely over the Internet. With remote management, Online Tech can provide the “hands and eyes” to help you manage your equipment as required at the data center.
- **Dedicated server** – is when Online Tech provides the server or servers that deliver the level of performance required to meet your performance needs. In so doing, you purchase a service level versus a capital asset and are no longer locked in to a owning or managing a specific version of hardware and software technology. Online Tech performs the hardware maintenance, required upgrades, patches, and responds immediately to all server hardware failures. Prior to installation patches or upgrades, Online Tech staff communicates with your application and database teams to test and assure proper implementation.

B. SERVER ENVIRONMENT AND MANAGED SERVICES

Online Tech goes beyond simply renting space in a data center or renting a server to our customers. We realize that the IDS requires more than space and servers to be successful. Online Tech delivers all the components and services necessary as a single source provider. Some businesses may only need a secure, environmentally controlled data center and filtered power, while providing the rest of their own environment. Other businesses choose to outsource everything including managed services, such as firewall, load balancing, and backup while others want to outsource all the way up to the web and e-mail services.

At Online Tech, our dedicated servers come bundled with the total environment essential to deliver the necessary level of security, reliability and scalability. In the case of co-location, customers have the option to select from our variety of managed services to fill any missing gaps in their IDS.

C. SERVER ACCESS

The physical access to the servers hosting your application or website can be classified in three levels: no access, business hours, and 24 x 7 x 365. Hosting vendors or infrastructure providers may provide one or up to all of these options. At Online Tech, we provide the security and flexibility to select any of these three access options to meet your pressing business needs.

D. SERVER STORAGE

Within data centers, servers can be placed in open area or in private locked cabinets depending on the level of security required. Most data centers provide either open ward or else locked cabinets. At Online Tech, you can have the flexibility to select any of the available options and change between storage options as their business needs evolve.

- **Open ward** – servers are placed within the secure walls of the data center and placed on open shelves or open racks. Though within the secure data center and cameras monitoring the area, it is the less secure option due to other open ward customers having access to the servers.
- **Shared locked cabinet** – the next level of server storage where your equipment is placed within a locked cabinet that is shared by one or more customers. The shared locked cabinet provides increased security by limiting the access to only the customers' sharing the cabinet space.
- **Private, locked (full or partial) cabinet** - customers needing increased security can select having their own locked space that no one else has access.
- **Private cage area** – customers requiring significant space for storing servers and private access can select private areas secured by a locked fenced cage.

Network
Internet Connection

VI. Network and Connectivity

A. NETWORK

The network architecture should be designed to reduce the threat of outside attacks. Online Tech uses numerous schemes to protect our network. Some these include: IP block assignment (separate subnet for each customer), redundant firewalls that restrict ping, notifications from cert, updates of all security patches as released, regular updates of firewall rules, intrusion detection systems, anomaly detection systems, wireless-free network, virus scanning on all incoming and outgoing traffic and regular review of network traffic logs and statistics.

Specifically, we use an Anomaly and Intrusion Detection system (next generation online security technology) that relies on the Snort technology. Firewall protection is provided by Cisco Pix devices and dedicated redundant Linux servers running IP tables. This allows us to provide virtual firewall services to each customer independently. All firewall hardware is redundant and backed up either by on-site spare parts or 7x24 contracts with 4-hour response time.

B. CONNECTIVITY

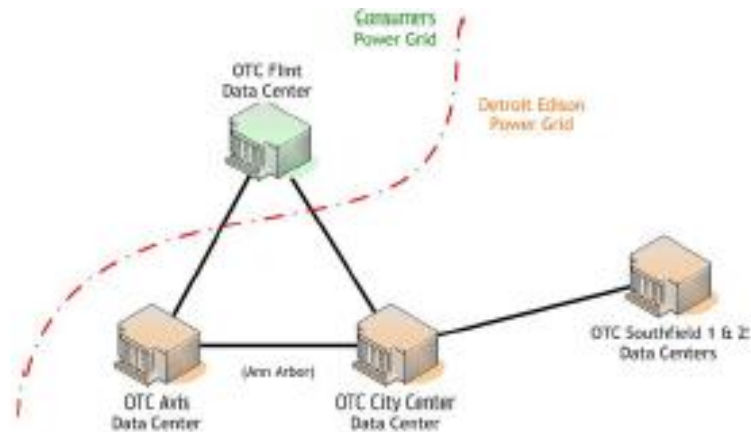
Some data centers also house CLECs or wholesale network providers that can provide high speed internet access. Online Tech's data centers are home for or closely located to a redundant set of CLECs and wholesale network providers each with their own fiber ring throughout Michigan and various regions of the Midwest. Because high availability network connections are as important as uninterrupted power to your internet application, all Online Tech data centers have two separate Internet backbone connections to two separate network providers. This redundant network connection through redundant providers maximizes the service level assurance we can deliver on network uptime.

We also allow our customers to use a pre-established connectivity provider, if they so choose. In this case, the customer is supplied a cross connect in the data center to their preferred provider.

VII. Data Centers

Electricity
Facility

Online Tech. has 3 High Availability Data Centers across southeast Michigan. These data centers sit across 2 different power grids with back-up power and redundant high bandwidth network providers to deliver a high availability, always-on computing infrastructure for our clients.



A. CENTRAL ANN ARBOR DATA CENTER FACILITY AND POWER

Online Tech's corporate office and it's central data center are housed in a seven story office building in downtown Ann Arbor, Michigan. Across the street from the police headquarters and the city's main fire station the building is ideally located for rapid response by Ann Arbor's emergency services.

This building is home to a number of CLECs and wholesale network providers each with their own fiber ring throughout Michigan and various regions of the Midwest. Across the street is the SBC main POP for Ann Arbor. The infrastructure in the building and to the data center can support redundant OC-48 connections.

The data center itself is on the 3rd floor which is approximately 30 feet off ground level. The chance of a breach through an outside facing wall is extremely low given the height from ground level. The remaining portion of the 3rd floor is occupied by Online Tech's offices.

The data center has six separate rooms: Battery and Power; Network infrastructure; Customer work room with KVM access; spare parts closet; NOC technician offices, helping to further segment risk and damage.

B. PRIMARY DISASTER RECOVERY DATA CENTER FACILITY AND POWER

Online Tech has a premier data center and backup / disaster recovery site located in Flint, Michigan. The facility totals 32,500 SF, which includes 22,000 SF of 24" Telco grade raised floor area and 8,000 SF of finished and furnished power-protected office space.



This world class facility was built by EDS for General Motor's, GMAC's, and Delphi's Disaster Recovery. The data center boasts 24" Telco grade raised floors, dual 8,000 volt, 2,000 amp power supplied via diverse paths from two substations. The power feeds are backed up with dual 1.3 Megawatt generators all fed through 1.6 Megawatt of synchronized UPS modules, 330 tons of cooling, biometric security, on AT&T's fiber SONET ring.

This data center offers world class capabilities for not only enterprise firms and service providers, but also individual businesses. Space in the facility can be acquired in large amounts as a production data center or disaster recovery site or per U as a colocation site. The facility is also equipped to meet the needs of those service providers and firms requiring office space for technical support staff.

C. SOUTH ANN ARBOR DATA CENTER FACILITY AND POWER

Online Tech's newest data center is a 10,500 facility previously used by Thomson/Medstat located in Avis Farms, a prestigious high tech park south of Ann Arbor. The data center is in a stand alone single story building. Biometric, key card and door locks provide physical security backed up by video surveillance.

The power infrastructure boasts separate natural gas and diesel generators delivering 850 KVA capacity. Power is conditioned by dual Liebert UPS with 375 KVA capacity. Cooling is delivered by 100 Tons of Liebert air handlers. The data center has dual fiber circuits connecting to Online Tech's other data centers.

VIII. Conclusion

Internet applications require many pieces to work properly to deliver a successful experience to the end-user. The Internet Delivery Stack is a useful framework for planning and tracking the myriad of components necessary to deliver successful Internet applications.

Internet-based services are expected to be *secure*, and *scalable*. Online Tech's Internet Delivery Stack can be used to describe each layer of components required to deliver a successful Internet application. The Internet Delivery Stack can also be used as a road map to develop High Availability applications and a Disaster Prevention and Recovery Plan (DPRP) for your critical data and applications so that management can make intelligent trade offs given their specific application and business state.

Online Tech understands that each business has requirements for their Internet based application and offers a broad range of solutions to meet even the most demanding needs. Call us and ask us about how we can help design your High Availability Internet Delivery Stack.