Ensim ServerXchange: the Key to Hosting Automation

A White Paper

July 2001

This white paper describes how Ensim ServerXchange allows you to grow your customer base, offer new hosted plans, and reduce costs through automation. We first outline the hosting life cycle -- the stages that you go through in dealing with customers, servers, applications, and services. We then consider the objects that you need to manage at each stage of the life cycle and lay out the set of tasks you need to handle to keep your business running.

Ensim ServerXchange provides an industrial-strength hosting automation solution for the majority of these tasks. We describe both the overall architecture as well as the patent-pending technologies underlying this product

1. The Hosting Life Cycle

A good way to think about the hosting business is to understand how the interaction between you and your customer changes over time. This usually follows a four-step *hosting lifecycle* consisting of service definition, provisioning, operations, and growth (see Figure 1).

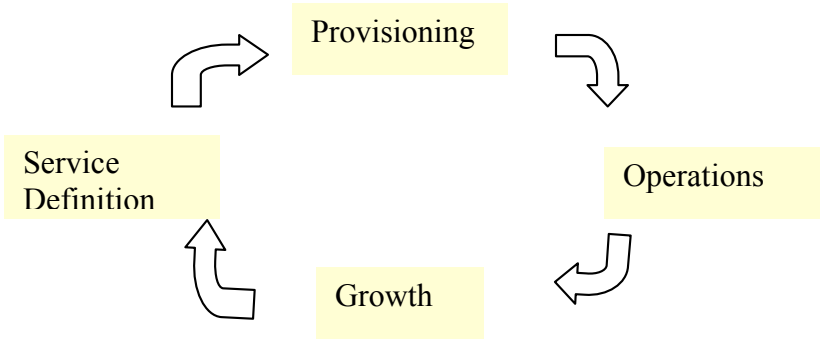

Figure 1: The Hosting life cycle

- The first step in the cycle is *service definition* where you decide what service plans to offer and at what price. A service plan bundles together a set of applications like Apache, FrontPage server, and PHP with a bandwidth quota, a disk quota, and other resource quotas.

- *Provisioning* refers to the steps required to actually deliver the applications and resource quotas defined in a service plan to a customer.
- *Operations* covers the day-to-day activities required to keep services available, customers satisfied, and money rolling in.
- *Growth* refers to the steps taken to deal with enhanced customer requirements, either for increase in resource quotas or for more services.

Let's look at these steps in more detail.

1. Define Service Plans

You typically offer a set of *service plans* to your customers. A service plan bundles a set of *applications*, such as email, FTP, and Web, running on a set of OS platforms, such as Unix, Linux, and Windows 2000, with a set of *resource quotas*, such as the amount of disk space allowed or the number of bytes that can be transferred every month. Your customers choose one or more plans depending on their needs. In this phase you define the set of plans you want to support and what you are going to charge for them.

2. Provision Plans

When a plan is selected, you have to *provision* the plan, that is, set up applications and servers in order to provide the promised services. For instance, if a customer requests a plan that provides a Web site, with a 200 MB disk quota, you have to, among other things, select a server with adequate capacity, create a directory for the customer's Web content, create a username and password for the customer to access the directory, and also make sure that the customer cannot exceed the storage quota.

3. Support Operations

Once the plan is provisioned, the customer can start using your services. Most customers will be happy with the default application settings. However, with complex services, a customer may want to make changes to one or the other feature. For instance, with email service, a customer may want to set up spam filters, autoresponders, vacation messages, or mailing lists. Each of these may require some sort of intervention by your support staff, typically fielding a phone call or email message, followed by several steps to implement the request.

At the end of the month or quarter, customers need to be billed. They may be eligible for volume discounts, or may have coupons they can redeem for service. They may need to be reminded about a lack of payment, or eventually have their service cut off.

Another set of operational issues arises from hardware or software failures. If a server fails, the applications hosted on that server become unavailable, hurting customers. You have to restore customer sites on another server. Similarly, if there is an application failure, the failure has to be detected and fixed. Finally, every application periodically issues upgrades. On receiving an upgrade, you have to convince yourself of the stability

of the new application version, then upgrade tens or hundreds of application instances – a challenging task! All of these tasks fall under the category of "operations".

4. Manage Growth

Finally, successful customers need more services, more capacity from their existing services, or more customization. You need to satisfy your customers by providing an adequate path for their growth.

2. What are the challenges?

Another way to look at the day-to-day tasks of a hosting provider is to consider the objects that need to be managed. At each stage in the hosting lifecycle you have to manage four things: servers, applications, customers, and infrastructure.

- *Server Management* deals with servers and operating system, independent of an application, such as creating user accounts, upgrading the server OS, backing up and restoring file systems, and setting user permissions and quotas.

- *Application Management* deals with application-specific tasks, such as installation, configuration, upgrade, and deletion.

- *Customer Management* deals with customers, such as registering a new customer, storing customer billing information, generating bills, following up on invoices, and tracking trouble tickets.

- *Infrastructure Management* deals with data center infrastructure, such as load balancers, firewalls, storage devices, and authentication servers.

If you combine the hosting lifecycle with the set of things you have to manage, you can easily figure out the set of tasks that you need to carry out to deliver hosted applications. This is shown in Table 1.

	Server Management	**Application Management**	**Customer Management**	**Infrastructure Management**

	Server Management	Application Management	Customer Management	Infrastructure Management
Definition	Decide server configuration(s) and operating system for each offered plan If virtual server hosting is supported, define virtual server configuration	Decide which applications will be supported in each plan Obtain application licenses from ISV if necessary	Design customer billing record format	Order and install all necessary infrastructure Establish relationship with Internet name registrar Establish relationship with credit-card verifier
Provisioning	Install server if necessary Select server and install OS on the server, if necessary Create user account(s) Set user quotas and permissions	Install applications corresponding to selected plan Set default parameters for selected applications	Create user record and fill billing information Send a welcome message to user with login information Give customer access to customer-specific Web information	Register domain name Obtain an IP address if necessary Set up load balancer and firewall if necessary
Operations	Backup server file system Detect and respond to server failure Upgrade server operating system	Detect and respond to failure Customize application to meet customer needs Upgrade application	Bill for application and infrastructure usage Record and follow up on customer complaints	Detect and respond to failures Customize infrastructure to meet customer needs
Growth	Increase server capacity if needed Move customer(s) to a more powerful server if needed	Obtain additional application licenses if needed Obtain and host additional applications and suites	Build in capability to deal with large number of customers Provide single sign-on to all services	Provide additional infrastructure if necessary

Table 1: Hosting tasks

From the table, it should be clear that hosting applications is a complex process that requires you to have a thorough understanding of applications, servers, and customers. Now, imagine that you get a hundred new customers deployed a day! Without a hosting automation solution, your only option is to employ large numbers of highly trained personnel. This, however, can dramatically reduce your return on investment. Moreover, there would inevitably be human error, missed steps, and other problems. To grow a

hosting business profitably, you *have* to automate the tasks identified in Table 1. Ensim ServerXchange is the industry leader in automating these tasks.

3. ServerXchange architecture: the physical view

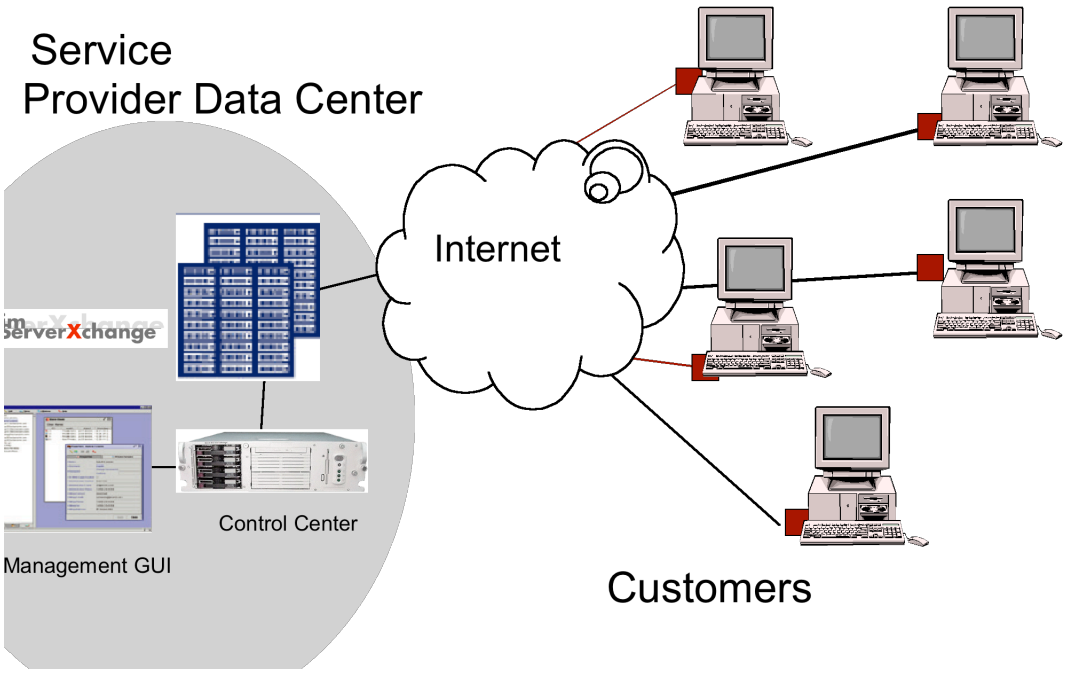

Figure 2: ServerXchange physical architecture. The Control Center communicates with agents running on each managed server and provides a single point of management and control. The management GUI communicates with the Control Center.

Figure 2 shows the physical layout of the ServerXchange product. ServerXchange software is shipped as part of a management appliance called a *Control Center*. You rack this appliance, power it up, give it an IP address, and can then use it to manage your infrastructure. You do so by pointing a web browser to the Control Center's IP address. This brings up a web-based user interface. The interface allows you to configure the Control Center, and manage servers, applications, domains, and customers.

To manage a server, you give the Control Center the IP address of the server and its root password. The Control Center loads a management agent on the server. This agent can then do things like reboot the server, monitor usage of the server, and so on. An icon representing the server shows up in the Control Center's user interface that provides a single point of management and control for that server.

The Control Center uses the management agent to distribute, update, and configure applications on a managed server. Icons representing managed applications appear in the Control Center's user interface to provide a single point of management for that application. You can use this icon, for instance, to update an application version on one or more servers, load a new application on a server, or configure application properties.

The Control Center exports a command line interface that allows you to easily integrate it into your existing infrastructure. Moreover, configuration information present in the Control Center database can be accessed using an SQL query or ODBC. This allows ServerXchange-managed infrastructure to be billed or viewed by another management platform.

Note that Ensim does not provide the managed servers nor, in most cases, the hosted applications. Service providers are free to manage any server running a supported operating system. Similarly, service providers are free to run any application on the manages servers. Ensim does make it easy to deploy and manage certain *wrapped* applications. We describe the notion of wrapping and the underlying technology later in this document.

4. ServerXchange architecture: the logical view

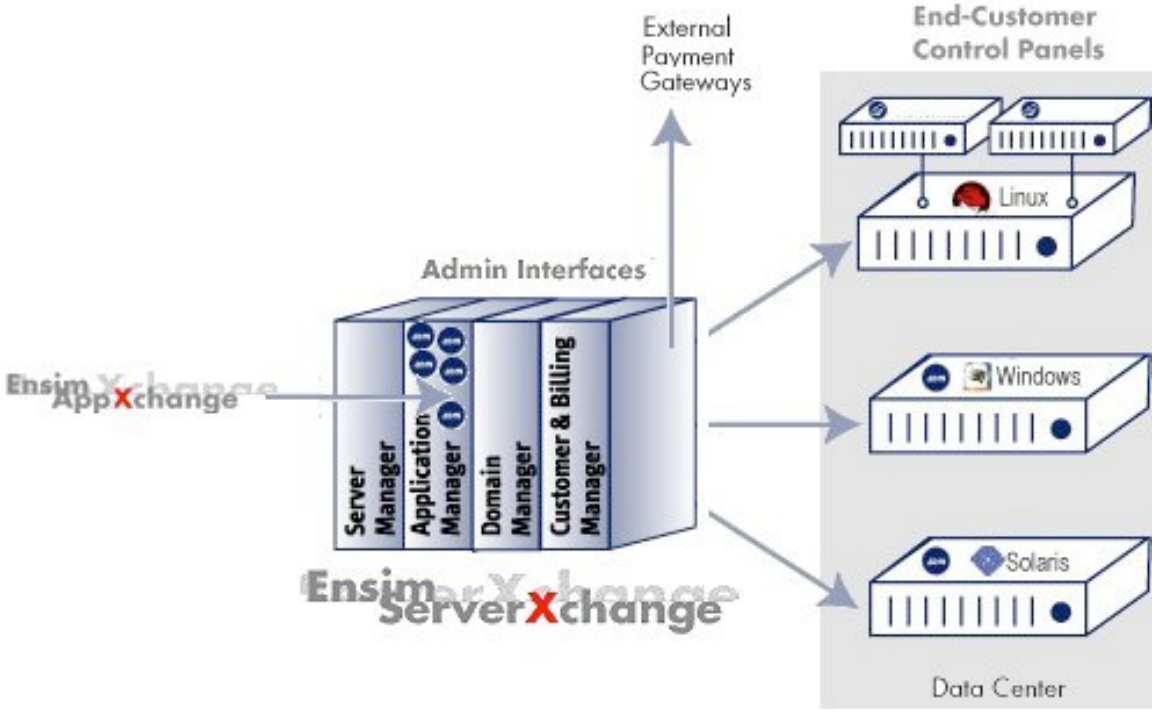

Figure 3: ServerXchange logical structure. The Control Center provides server, application, domain, and customer management to Solaris, Windows, and Linux servers. An ISP's customers access their servers and services using Web-based control panels.

Figure 3 shows the logical layout of the various ServerXchange components. The Control Center management appliance provides four types of functionality: server management, application management, domain management, and customer management. You can access this functionality using the ISP interfaces and your customers can self-manage their applications and servers using the customer interfaces. The Control Center also connects to the Ensim AppXchange service to receive application updates, new

applications, and ServerXchange updates. In the rest of this white paper, we discuss each management function at length.

5. Server management

ServerXchange provides a comprehensive set of tools for server management. At a basic level, it allows you to load an operating system on a server, monitor its status, and reboot it. At a more advanced level, it partitions a single physical server into a number of isolated *private servers* that each receives its own guaranteed share of server resources. Private servers can be resized (that is, given a different resource share) and moved from one physical server to another. We now discuss basic and advanced server management in greater detail.

5.1 Loading an operating system on a server

A ServerXchange Control Center serves as a repository for operating system images. These are either pre-loaded (for free operating systems) or uploaded by you using a Web-based interface (for commercial operating systems). Once an OS is available on a Control Center, it can be loaded on any managed server. The idea is that most commercial servers come pre-loaded with an operating system. However, this OS may not be the right version, have the right patches, or be of the right kind. In order to create a single homogeneous OS image on its servers, you can load a standard OS image into the Control Center, then use it to download it to every managed server. The Control Center does so by creating a ‘kickstart’ floppy disk. You need to simply insert this floppy disk into a server and reboot it. When the server comes up, it will load the right OS type and version.

5.2 Server monitoring

All the servers and private servers at your data center can be monitored from a single console, shown in the figure below. An icon in the Control Center user interface represents each server. The Control Center periodically pings the server. If a server does not respond to repeated pings, then it is presumed dead, and the Control Center icon is

suitably updated.

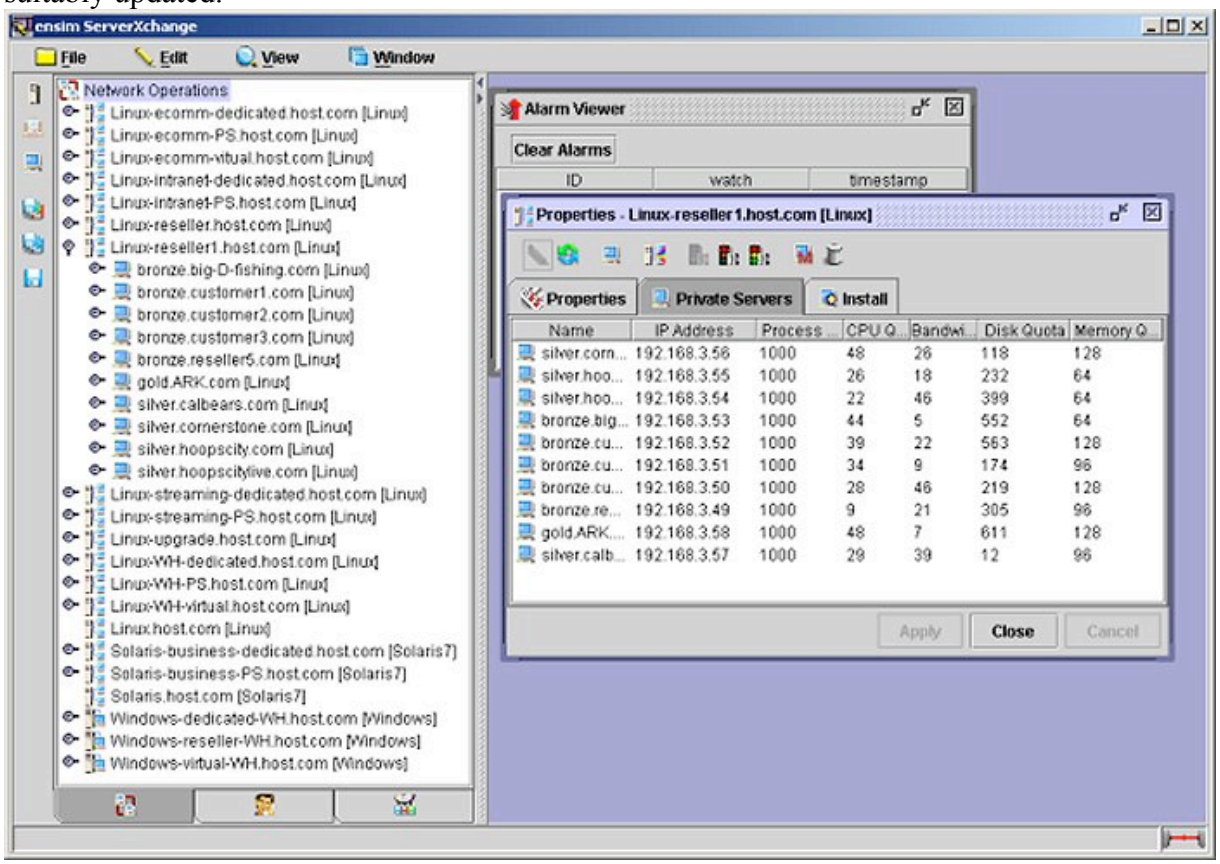

5.3 Private Servers

ServerXchange lets you partition a server into a number of *private servers*. A private server essentially is a set of processes that are associated with a resource context. Specifically, every resource access from a process within a private sever is modified so that:

- Resources outside the scope of the private server are not visible
- The amount of resources consumed by a private server are limited by a quota

The net effect is that to a customer, each private server appears more or less like a physical server. They can telnet to it, compile any executable, run any binary, or access most OS resources, such as files, IPC keys, /proc, or mount points, without knowing that they are sharing a single physical server with other private servers. A private server allows a physical server to be shared among multiple customers without their knowing about each other's existence and without their being hurt by excessive resource consumption by others. In particular, private server owners can be given root access, yet they cannot access other private servers, nor use resources beyond those granted to them by the ISP.

A private server is guaranteed four types of isolation:

- **Fault isolation**: if a private server crashes, it does not hurt others. A private server can be rebooted without affecting any of the other private servers sharing the same physical server.
- **Functional isolation**: each private server is given a functionally separate view of the physical server. For instance, each private server has its own file system, its own IPC keys, and its own view of the network connections.
- **Performance isolation**: each private server is guaranteed its own share of physical sever resources such as CPU, memory, disk, and network bandwidth. So, a private server that tries to consume all the CPU in the physical server, for instance, will not be able to use more than its fair share.

  Private servers have two types of bounds: *max* bounds and *min* bounds. A max bound restricts a private server to a maximum share of a resource. Max bounds are used for disk quota, virtual memory size (i.e. the sum of physical RAM and swap size) and the number of simultaneous telnet sessions supported by a private server. Min bounds guarantee a minimum share of a resource independent of the load on the system. With a min bound, if a private server can potentially obtain more than this minimum share without hurting others, it is allowed to exceed the bound. So, for instance, if only a single private server is active, it can get a 100% resource share for all resources that have min bounds. Private servers support min bounds on CPU consumption and network bandwidth.

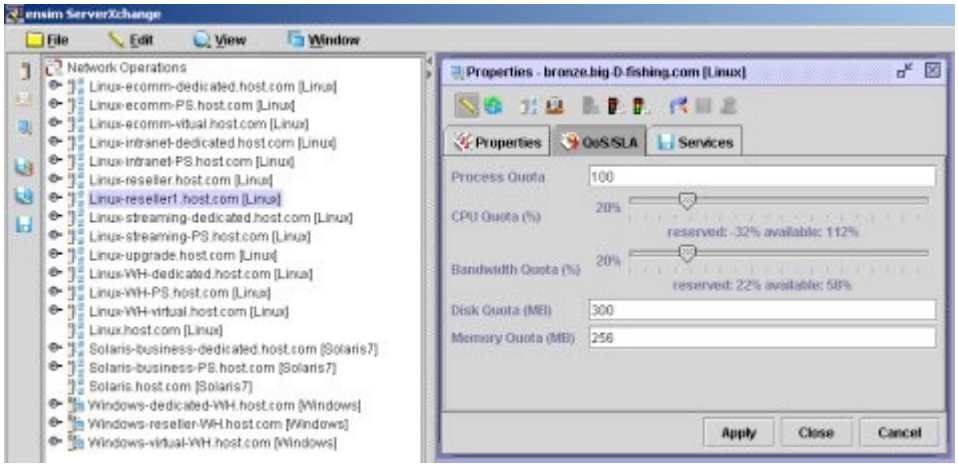

- **Address isolation**: each private server has its own IP address, and can only receive packets destined to this IP address. Moreover, applications on a private server are not allowed to spoof the source IP address, greatly enhancing security.

A private server allows you to offer two new types of services: true reseller access, and virtual dedicated servers. You can partition a server into private servers, and give each reseller full access to the private server without hurting other resellers sharing that server. This greatly reduces the cost of reseller hosting while still giving a reseller complete control over their operating environment. In conjunction with the WEBppliance Web Hosting Appliance software, and the Reseller Automation Suite for customer billing, Ensim ServerXchange provides an unparalleled set of tools for managing a reseller channel.

A private server also allows you to offer a 'virtual dedicated server'. This is just like a dedicated server, except that a customer receives only a share of the physical server, instead of its entirety. Note that a private server differs from a physical server in three ways:

- A private server cannot access raw memory (/dev/kmem in Linux)
- A private server cannot access raw disk (inodes) and most physical devices (for instance: mouse, keyboard).
- A private server cannot insert and delete kernel modules

In practice, we find that very few applications are affected by these restrictions. In fact, a non-root user can distinguish between a physical server and a private server only with great difficulty. Other than these differences, a private server supports all the APIs supported by a standard Linux or Solaris operating system.

5.3.1 Private Server overhead

A private server virtualizes the underlying operating system, and like any virtualization scheme, incurs some overhead. This overhead depends heavily on the workload. We measured private server overhead for web hosting using SpecWeb99. The results are shown in the figure below:

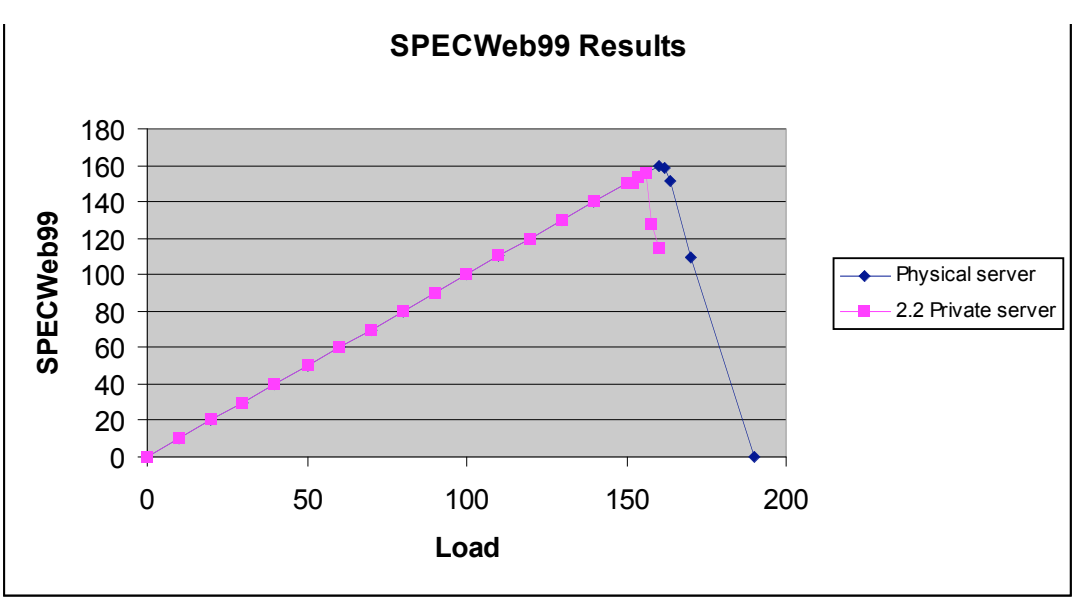

The X axis shows the load, measured in number of connections, placed on the web server. Each connection from a client is at 400kbps: a load of 150 connections, therefore, corresponds to 60Mbps of sustained load. The Y axis shows the SPECWeb result. This is the number of connections that receive a throughput of at least 320kbps.

With a web server running on a physical server, the highest SPECWeb number achieved is 162. If the load exceeds 160 connections, the SPECWeb number drops rapidly due to heavy swapping. To determine the overhead due to a private server, the same web server is run on an identically configured private server. On a private server, with the same configuration, a SPECWeb of 155 is attained: an overhead of 4.32%. Another way to measure overhead is to look at the response time degradation. At the peak sustainable connection rate of 150 connections (60 Mbps), we find that the response degrades from a mean of 336 ms to 343 ms, an overhead of 2.04%. These results are summarized below.

SPECWeb99 result with physical server	**SPECWeb99 result with private server**	**Overhead**
162	155	4.32%

Response time for physical server at peak load (ms)	**Response time for a private server at peak load (ms)**	**Overhead**
336	343	2.04%

Very similar results are obtained when comparing a Solaris physical server with a Solaris private server.

5.3.2 Differential QoS

ServerXchange has the ability to allocate different quality of service guarantees to different customers. To test this, we partitioned a server into four private servers. Two private servers were given a 30% share of network and CPU resources, and a memory limit of 128MB. Two other private servers were given a 15% share of network and

CPU resources, and a memory limit of 128MB. The two 'high-end' private servers, therefore, have twice the resource allocation of the two 'low-end' private servers. All four domains were loaded in parallel by four SPECWeb load generators. The table below summarizes the observed results:

Domain type	Operations/sec	Response time (ms)	Bitrate (bps)
Low 1	6	2887	42088
Low 2	7	2811	42585
High 1	14	1412	85686
High 2	14	1416	84780

Note that:

- Both low-end and high-end domains get the same performance in terms of operations/second, response time, and bitrate.
- High-end domains get twice the performance of low-end domains for all three metrics.

This demonstrates the ability of ServerXchange to provide differential quality of service to different private servers.

5.3.3 Per-managed server resource usage

When a server is brought under the ServerXchange management umbrella, the management agents take up about 360K of RAM. They also reserve a minimum of 10% of CPU and bandwidth for themselves (this doesn't mean they use all of it, it just means that if they need to, they can use at least 10% of the capacity. In practice, they use an average of less than 1%, but at times when you are adding a Private Server or adding a service, they may burst all the way to 10%). In addition, each private server uses disk space on the box to cache services and applications. If no services are installed, then we take 300 MB for the base system: additional services take up more disk space.

When a private server is added, each additional private server takes an additional 200 K of kernel memory. We use an additional 300 MB of disk space for the base file system (this comes out of the disk quota assigned to that PS). We do not tie up any additional CPU or network bandwidth, other than a small amount used to report PS statistics to the control center. When a private server is moved or removed, these resources are released and are reusable. These figures are summarized in the table below:

Resource	At physical server registration time	Per 2.2 Private Server
Memory	360K	200K
CPU	10% reserved	Chosen by operator
Network	10% reserved	Chosen by operator

bandwidth		
Disk	200MB	300MB

A capacity planning Technical Document, available to Ensim customers, gives more details on sizing physical servers to run ServerXchange.

5.4 Storage management

A server's storage can be either local or located on a file server, typically on a Network-Attached-Storage (NAS) server. Traditionally, a server uses local disk for storage. A hosting provider, however, may want to share file systems, particularly read-only directories, among multiple servers. This serves two purposes. First, the overall disk usage is reduced. Second, and more importantly, this ensures that all servers share the same version of the files.

A shared file-system solution can be achieved using NAS and the NFS protocol. In this solution, a NAS server holds read-only directories and exports them to the network. NFS clients, one on each server or private server, can then import a shared directory into any mount location on their local file system. With this arrangement, multiple servers of private servers can efficiently share libraries and binaries.

Ensim-managed servers can use both local and NAS storage. When used with local storage, multiple private servers can share a single disk volume. Each private server, however, is given the impression that it owns the entire volume and has its own base file system. When used with NAS storage, a private server can mount an NFS-exported directory into its local file system. Multiple private servers are allowed to mount the same NFS-exported directory.

5.5 Loose clustering for high availability

Some hosting providers like to move all but a handful of files from a server's local disk to a shared network server. This has both pros and cons. On the pro side, this means that the service provider has a single storage server to backup and restore. Moreover, if a compute server were to fail its state is captured in the storage server and is easily restored. On the con side, the shared server is a single point of failure. Performance, too, can be a problem if the network bandwidth is lacking, or if 'hot' shared files need to be locked often. Nevertheless, the manageability of shared storage has made this a very common feature at high-end hosting providers.

Ensim ServerXchange supports this style of storage management. When a private server is installed, its mount location can be either a local directory or an NFS-exported directory. With the latter option, all private server files are accessed over NFS. With this arrangement, if a compute server fails, a private server assigned to that compute server can still be moved to another compute server, where it will restart with the right IP address. ServerXchange, in conjunction with NAS storage, therefore allows hosting

providers to quickly recover from server failures, and provides a loose form of clustering with no need for an additional load-balancer.

5.6 Moving private servers

A private server can be moved from one physical server to another. A move consists of shutting down the private server, moving its identity and associated file system to another server, then restarting the private server. If the private server file system is on a NAS server, then a move can be accomplished in a matter of minutes. Move is also supported with local file systems, in which case the files are copied from one server to another. In this case, the time to move the private server depends on the size of the file system and the data center network bandwidth.

The 'move' capability of a private server uniquely decouples the service provided to a customer from the physical server providing the service. The only binding between a customer and the service is the private server, which can be moved between physical servers. Moreover, the private server can be given varying amounts of resource quotas. This allows a ISP to give a customer a small resource quota, then seamlessly upgrade the service to deal with customer growth. If a customer outgrows the resources of a dedicated server, service can be moved to a more powerful server with practically no customer downtime.

5.7 Security implications

Security has emerged as a major challenge to hosting providers. Web sites and hosted servers are constantly under attack. We can classify security management under two heads: attack prevention and attack remediation. Attack prevention consists of 'locking down' servers, that is, removing inessential services, applying security patches to applications, and, in general, following best security practices. Nevertheless, even the best-secured servers are liable to be compromised. One of the most troublesome problems in dealing with a compromised server is that an attacker can place a network interface card in promiscuous mode, thus snooping on every packet that traverses the network. A ServerXchange private server does not allow access to the network interface card. Moreover, it actively prevents packet spoofing – every packet leaving the private server must have the IP address of that private server. For these two reasons, a compromised private server is less harmful than a compromised dedicated server.

ServerXchange also eases recovery from attack by allowing a hosting provider to quickly set up a new server, installing an operating system and the hosted services. The ISP can then restore a customer's information from a backed up copy on to this restored server.

6. Application management

Ensim ServerXchange allows ISPs to manage mass-scale application distribution, configuration, update, and control. The basic unit of management is a *wrapped* application. An application wrapper, which can be written either by Ensim or by you,

provides a standard set of interfaces to install, start, stop, and update an application. A wrapper is similar in intent and format to the RedHat RPM, Sun pkg, or Microsoft MCSI formats. It allows automation scripts to manipulate applications in a uniform manner.

Application management follows a four stage pipeline. In the first step, Ensim or you identify an application that is both host-able and has wide-scale demand. Examples of such applications are Apache, majordomo, Microsoft IIS, and Allaire ColdFusion. Once an application is identified, the next step is to write a wrapper for it to make it manageable with ServerXchange. Most wrappers are straightforward to write, and can be written in a matter of hours or days. However, applications that have strict licensing requirements, or that modify the Windows registry in non-standard ways, need more complicated wrappers. Ensim provides a toolkit and detailed instructions to wrapper-developers.

Once an application has been wrapped, it is ready for distribution. In this second step you can prepare an application for distribution in two ways. First, wrapped applications can be uploaded into a Control Center and registered with it. This allows ISP-wrapped applications to be mass-deployed without further help from Ensim support. Second, an you can download wrapped applications from the Ensim AppXchange website. These downloaded wrapped applications need to be registered with the Control Center.

The Control Center displays all registered applications in the management GUI. It allows these applications to be installed on one or more managed servers or private servers. Given a wrapped application, installing it on a managed server is as easy as pointing and clicking in the GUI. ServerXchange eliminates tedious and error-prone techniques for application distribution that most ISPs use today.

Application versions change over time and some applications need to be patched to deal with security problems. A Control Center allows patches and updates to be distributed in the same way as an application, letting you mass-deploy a patch to all or a selected subset of its servers very easily. The screen shot below shows an example of mass update of an application across multiple private servers. Here, the Analog Web Log Analyzer 1.4 application can be updated on four private servers with a single click of the ‘Update’ button.

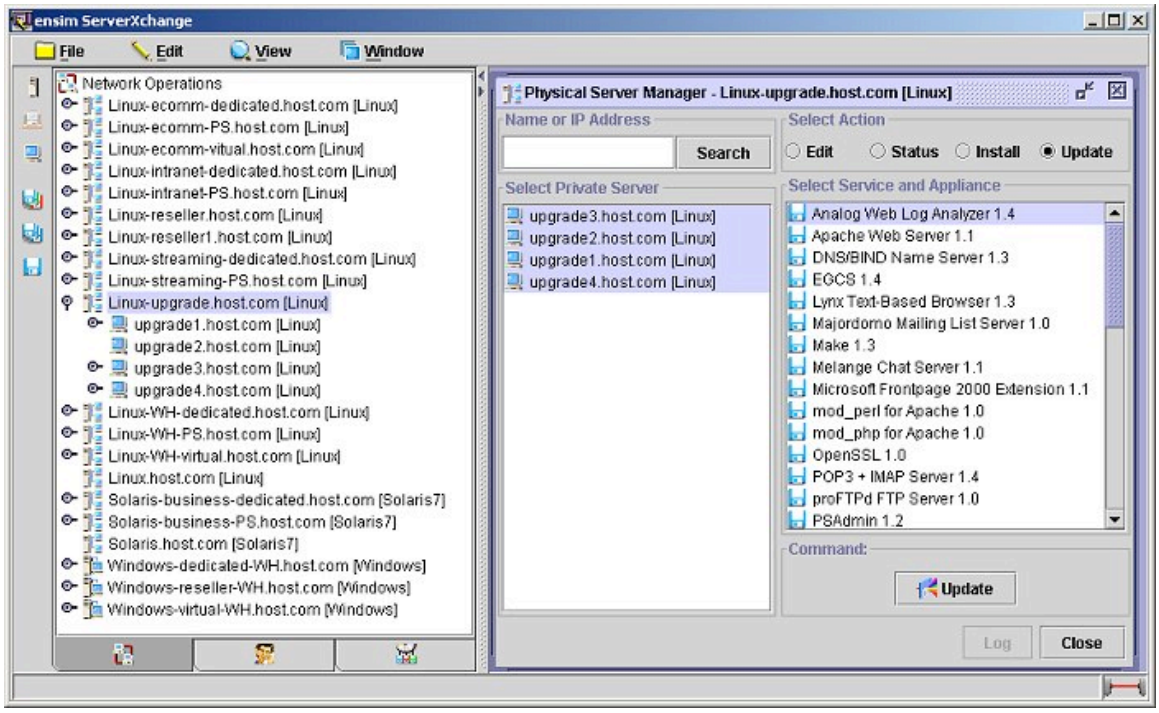

6.1 Application configuration and control

A hosting service provider ultimately is providing its customers with access to a hosted application. Almost every complex application supports tens or hundreds of configuration parameters and a significant challenge facing every provider is allowing their customers to manage the configuration of their hosted applications. The easiest way to manage this is to allow customers to call in or email their configuration requests. However, this method quickly uses up scarce support resources, preventing you from scaling your operations, and increasing costs. As a rule, ISPs therefore use Web-based *control panels* to allow customers to configure and control their applications. A control panel provides a way to authenticate customers, and then gives them a GUI to access application settings. With the right set of control panels, your customers can manage their own applications, freeing you to deal with other issues.

Ensim ServerXchange ships with a companion product, WEBppliance, that provides control panels for a dozen or so widely-hosted applications. These include Apache, sendmail, majordomo, ColdFusion etc. WEBppliance greatly increases your ability to scale while still providing their customers with highly-customized application configuration.

An important feature of WEBppliance is its notion of multiple levels of management and control. It provides application controls to an ‘Appliance Administrator’ who can install applications, create websites, and set top-level configuration parameters. At the next level, a ‘Site Administrator’ is allowed to change settings within a site and to create users who have access to a particular website. A Site Administrator can also change some application settings, such as those for email and Apache. At the lowest level, a ‘Site User’

is allowed to log into a website to upload or download files, or access their email. By providing hierarchical control to the suite of web hosting applications, WEBppliance allows you to devolve the responsibilities for managing web hosting to the right administrative level.

7. Domain management

Although a WEBppliance allows you to comprehensively manage the set of web hosting applications on a single server or private server, you may still need a way to coordinate domain creation and management across multiple WEBppliances. For instance, you may want to create a domain (also called a virtual site) on one of several WEBppliances, choosing the target WEBppliance from a single management console. Other functions, such as IP address range management, also naturally span multiple WEBppliances. The Ensim Domain Manager provides a single point of management and control across multiple WEBppliances. It allows you to manage Appliances, IP addresses, domains, and monitor the status of core services such as HTTP (web), FTP, SMTP (email), and POP, and monitor the total web usage data per server.

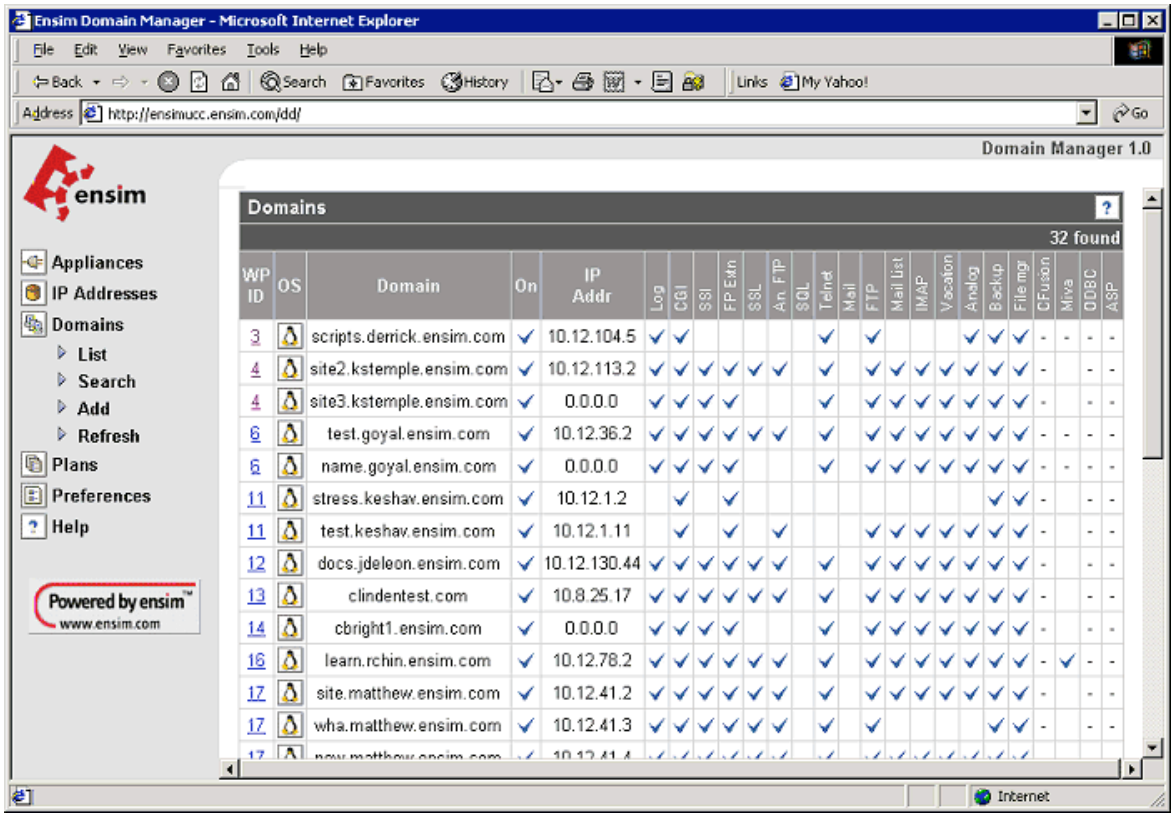

The Domain Manager displays a list of all appliances and domains present at a Control Center. From the appliance list, you can directly access the WEBppliance Appliance Administrator interface. You can select the customers to monitor when displaying

appliances. This allows you to hide WEBppliances being used by resellers rather than directly for virtual domain hosting.

You can use the Domain Manager to track and manage all the IP addresses being used for virtual domain hosting. The IP manager lists all IP addresses available and in use. IP addresses can be marked in use or out of use or can be deleted. When creating IP-based domains, IP addresses will automatically be chosen from a valid IP address range.

You can also synchronize their IP Address list with what is actually deployed. The synchronize feature polls for all the IP Addresses used by all the domains. IP Addresses not in the list are added and set as used. IP Addresses in the list and set as unused are set as used.

Finally, the Domain Manager allows you to create 'plans'. A plan is a standard set of features created along with a particular domain. For instance, a Bronze plan may create a domain that provides access to HTTP and FTP, 5 MB of disk, but no access to telnet. A Silver plan may add telnet access and increase the disk quota to 10 MB. The Domain Manager allows you to create an instance of a Bronze or Silver plan with a single click. This high degree of automation greatly eases the creation of shared hosting domains. It also eliminates human errors in creating virtual domains. You can create any number of plans, modify existing plan definitions, or change the plan associated with a particular domain.

8. Customer management

ServerXchange provides a comprehensive set of features for customer management. There are four steps in the interaction between you and your customer. Even before the customer purchases a service, you need to define a set of hosting plans, similar to

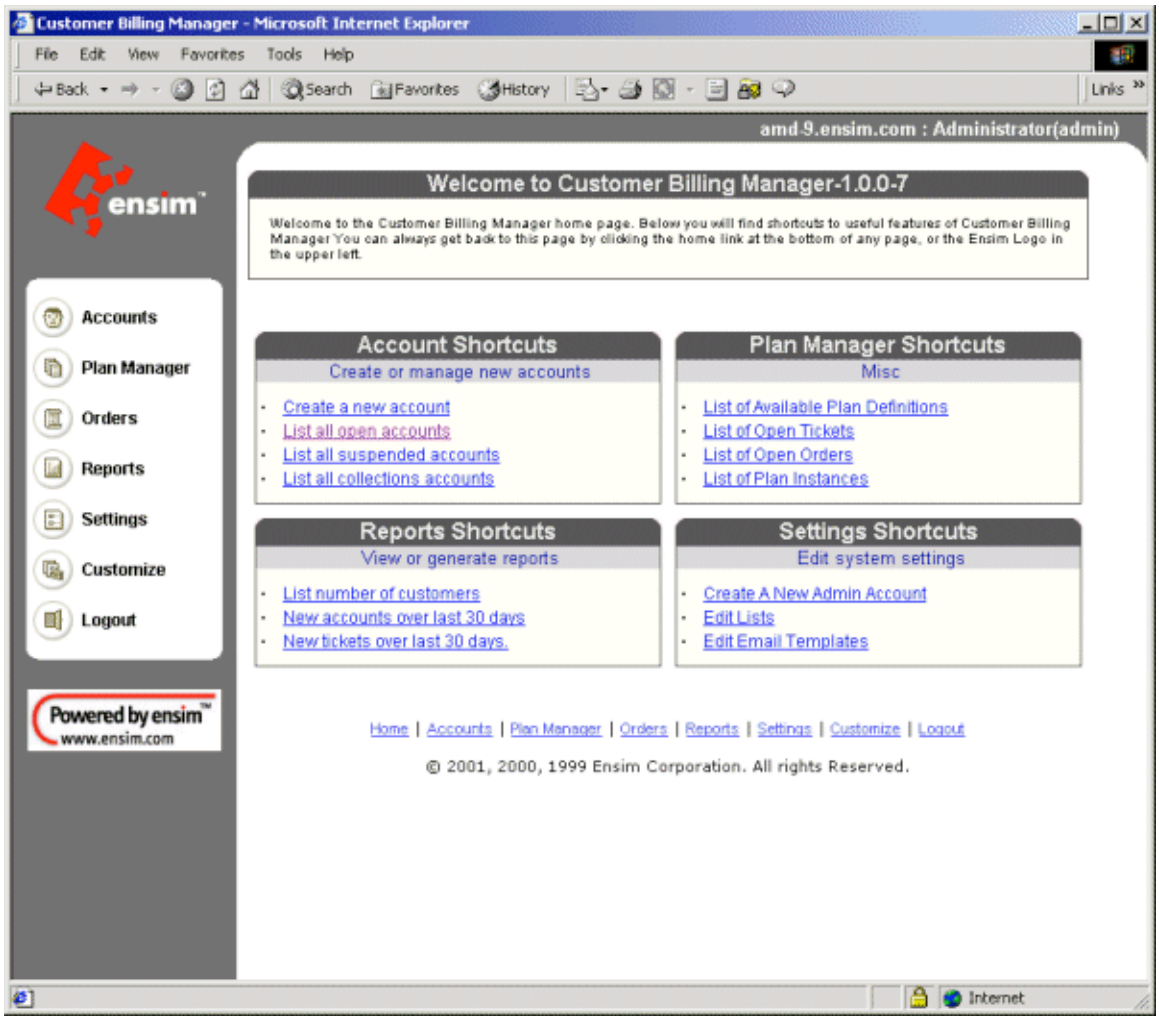

those described in Section 7. ServerXchange Customer and Billing Manager (CBM) allows you to define an arbitrary set of attributes into a plan. So you can use CBM to manage both Ensim and non-Ensim software, hardware and infrastructure.

When a customer wants to purchase a plan, you need to create an *account* for that customer. The account is associated with customer information such as a contact address, a credit card to bill, and so on. The CBM maintains a database of this information that can be accessed using an intuitive GUI.

Over time, a customer needs to be billed for plan usage. The CBM allows you to automate this process: customers are sent periodic email to remind them of payment due. If a customer pays using a credit card, then the CBM can automatically contact a payment gateway to obtain payment. Overdue accounts are automatically reminded of past due amounts, and after a certain number of reminders, user plans can be automatically placed on hold, or terminated. On another note, a customer may want to report problems to you. CBM provides a built-in trouble ticket system to accept and track trouble tickets.

Finally, some customers may purchase more plans, or may want to modify the set of attributes in their hosting plan. These can both be accomplished using the CBM GUI simply by associating an account with a different plan, or creating a customized plan for that account. CBM tracks the account history for every account, and also provides tools to statistically analyze the customer base, for instance, allowing you to track the proportion of overdue customer accounts over time.

In sum, the Customer and Billing Manager provides:

- A framework to define plans: their features, options, and billing parameters. Includes predefined templates for popular offerings.
- Automatic recurring billing and invoicing.
- Built in storefront generation, order administration, and tracking.
- Real-time credit card processing through popular payment gateways.
- Complete account history for every customer.
- Comprehensive reporting, trend analysis and profiling of customer base, for tracking and forecasting.
- Support for manual interventions such as crediting a customer, adding an extra charge, locking an overdue customer's account, etc.
- Customer interface for self-management, purchasing, reviewing transactions, and support.
- Trouble ticketing system.

9. Reseller support

Ensim ServerXchange provides a number of unique features to allow you to add on significant revenue streams from resellers. You can create reseller plans with a consistent look and feel on both Linux and Windows 2000. A combination of private servers and WEBppliance allows you to offer resellers:

- Root access to virtual dedicated servers
- Multiple levels of management on WEBppliance
- Comprehensive Customer and Billing Management for their customers
- Ability to resell applications like streaming and E-Commerce

Except for very small resellers, most resellers want to have root access to the server where their domains are hosted. This lets them customize their own settings and add or delete domains directly without contacting you. Usually this means having to dedicate a server per reseller, even if the reseller does not fully utilize it. Ensim private servers enables you to offer all the functionality that a reseller gets on a dedicated server - including root access, a guaranteed share of resources (including CPU, bandwidth and memory), isolation from other domains, and complete flexibility - all without dedicating a server per reseller. This gives you the ability to

- Lower your entry-level offering for resellers wanting root access.
- Create fully functional reseller plans with the ease-of-use of an appliance at tiered price points to target any size of reseller.
- Grow your plan with your reseller by simply migrating him up the tiers with a single drag-and-drop operation
- Increase your ROI per physical server

Reseller plans based on ServerXchange combine the ease-of-use and low support costs of an appliance with several unique features.

- Automated installation of services that enable the reseller to instantly set up and configure a domain without needing to contact you.
- Automated services management that enables the reseller to add and modify services on the fly with simple drag-and-drop operations.
- Control panels at three levels: appliance administrator for the reseller, site administrator for the domain owner, and user administrator for the domain user

One of the biggest headaches resellers face is figuring out billing details for each individual customer every month. This, along with customer support, is where they end up spending a large chunk of their time, when they would much rather spend their time on design, relationship building, marketing, or other activities. The Ensim platform enables you to offer your resellers a complete automation solution that automates billing and trouble-ticket management. This is similar to the Customer and Billing Module described in Section 8, except that it is now available to each of your resellers.

The Reseller Automation Solution includes

- Automatic recurring billing and invoicing for the reseller's customer.
- Automatic customer sign-up from the resellers website
- Automatic provisioning of plans for the reseller's customers .
- Centralized Domain, IP Address and Appliance Management for the Reseller.

- Framework for the reseller to define plans, with predefined templates for popular offerings.
- Built-in storefront generation, order processing and tracking for the reseller.
- Real-time credit card processing through popular payment gateways.
- Reporting, trend analysis, profiling of reseller's client base.
- Customizable customer interface for self-management, purchasing, reviewing transactions, and support for reseller's customers.
- Trouble-ticketing system for reseller.

ServerXchange allows you to host resellers for the following applications:

- Microsoft Exchange

  Resellers can offer their small business customers with messaging and collaboration solutions based on Microsoft® Exchange. The Ensim Exchange Appliance enables this by providing the capability to share a single Exchange server across multiple organization in a simple and reliable manner.

- Miva Merchant E-Commerce Plans

  Resellers can easily offer E-Commerce functionality, based on the popular Miva E-Commerce software, to their clients.

- Real and Windows® Media based streaming media

  The Ensim Media Appliance enables resellers to offer their customers media streaming capabilities based on RealSystem Server and Windows® Media Server, by sharing streams across domains.

10. Ensim Control Panels

Ensim ServerXchange delivers a set of comprehensive Control Panels for resellers, site owners, and site users that enable the simple service management of servers, Private Servers, and services. Ensim provides one set of Control Panels to manage Ensim WEBppliance and another to manage Private Servers.

Within WEBppliance, the Appliance Administrator Control Panel provides the appliance owner, generally a reseller, with a simple interface to manage popular Web hosting software such as Apache, DNS, Mail, and FTP. It allows the addition of IP and name-based sites, as well as the rapid creation and deployment of hosting plans via a template creator, where DNS, Mail, FTP, FrontPage, backup/restore, Web logs, file manager, and many other applications can be enabled and disabled. It also comes with a customizable

navigation bar, so that custom CGI scripts may be introduced into the hosting environment.

The Site Administrator enables site owners to perform routine administration tasks, such as obtaining access to many critical management functions, including Web, mail, mailing list, FTP, and database managers. The Site Administrator also provides the site owner with user manager, file manager, site usage reports, and backup and restore capabilities.

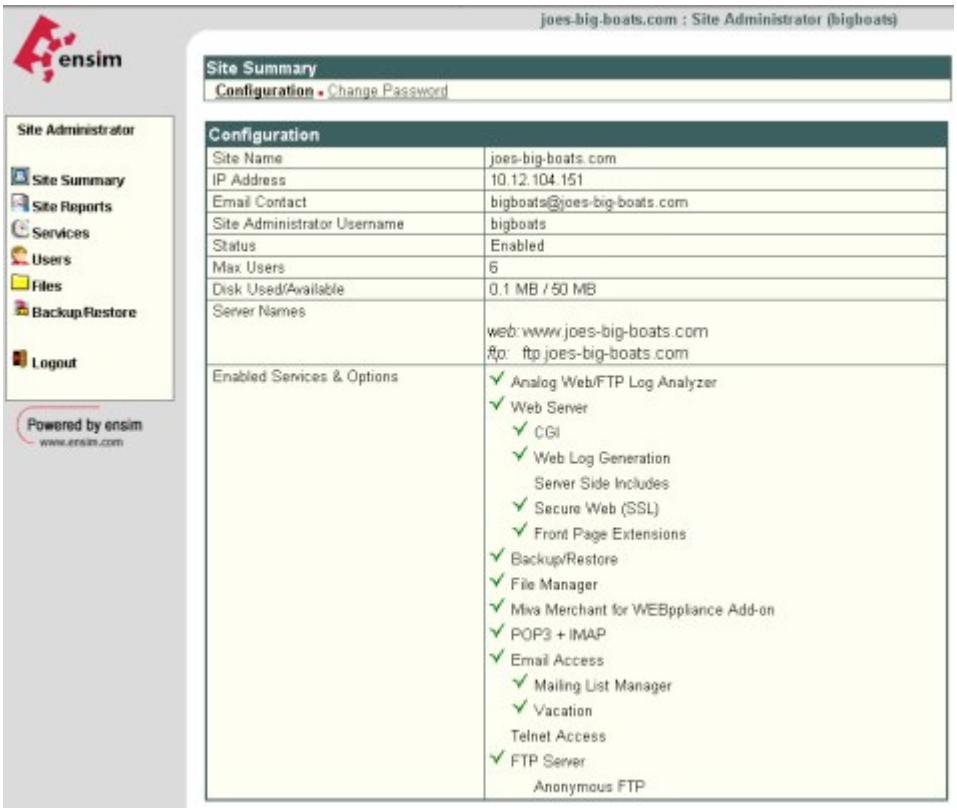

User Administrator gives site users the ability to manage their individual accounts. Using this Control Panel, users are able to access management functions such as user and mail managers, usage data, and backup and restore capabilities.

Ensim provides an additional set of Control Panels for the administration of individual Private Servers. The Account Manager interface is used by the owner of one or several Private Servers to view account information, Private Servers, and services under control, view information on individual Private Servers and customers, and add and update applications.

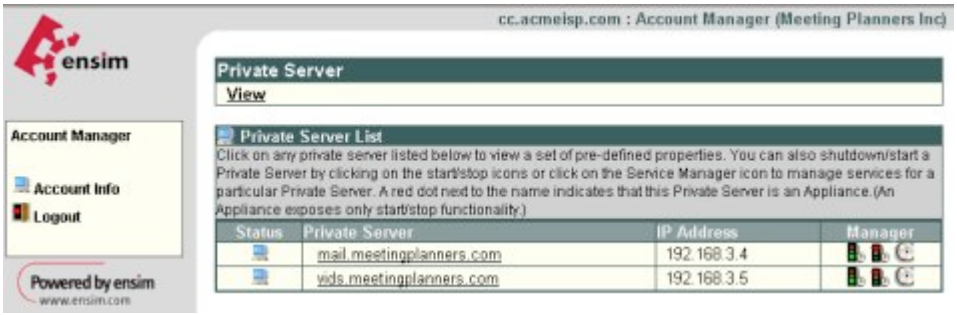

Administrators of individual Private Servers use the Private Server Administrator Control Panel to view network settings, system time, users, and running processes; view Private Server services, users, and files; manage application services and create users and e-mail accounts; manage files and file systems.

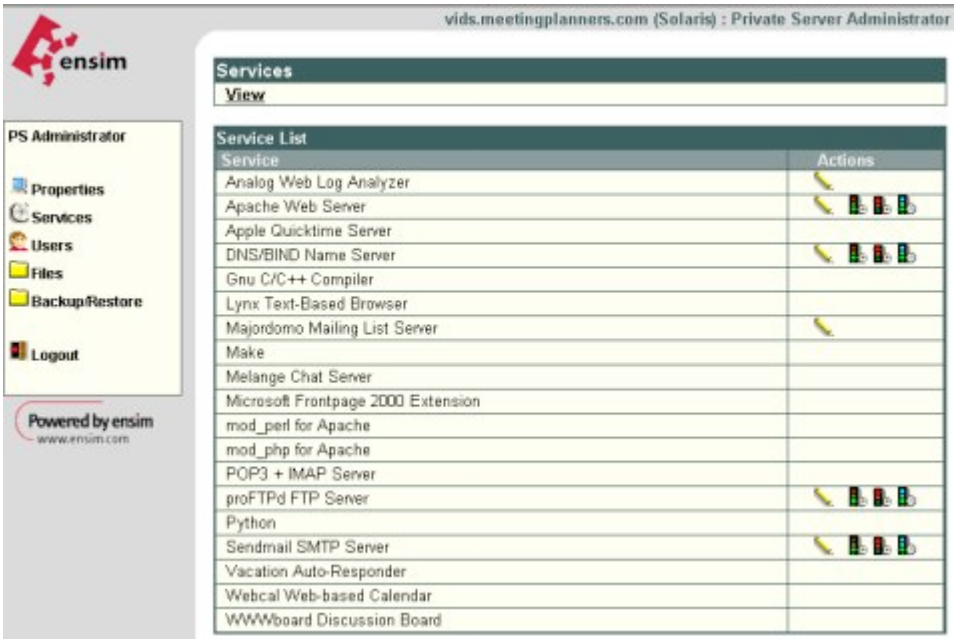

Users of individual Private Servers use the Private Server User Administrator Control Panel to upload and download files; set file and directory permissions, and view user properties and files.

11. Conclusion

1. *Completeness*

A hosting automation solution should automate as much of the hosting life cycle as possible, providing solutions for server management, application management, customer management, and infrastructure management. Moreover, the solution should integrate the various hosting tasks. For example, a system administrator should be able to pull up not only a customer billing record, but the set of servers on which customer applications are hosted, their current status and configuration, and any pending customer requests.

Ensim ServerXchange has the most-complete hosting automation solution in the market today. Its comprehensive support for server management, application management, domain management, customer management, and reseller management, along with a large set of control panels make it possible to automate virtually every part of your operations.

2. *Customer self-management*

To scale, you have to delegate as much of the hosting task as possible to end customers. These end customers fall into one of three categories: site creators, site administrators, and site users. A hosting automation solution should provide a set of *control panels* that allow each of these customers in self-managing their service. A well-designed set of control panels can greatly reduce your support burden.

As described in Section 10 Ensim ServerXchange supports comprehensive multi-level control panels that allow you to delegate almost all of the day-to-day hosting tasks to your resellers and customers.

3. *Support for a reseller channel*

Many ISPs have discovered that they can greatly extend their reach and revenue stream through resellers. A hosting automation solution should allow you to delegate control to resellers: the better the solution supports resellers, the easier it is for you to gain revenue from this rapidly growing channel.

As described in Section 9, Ensim ServerXchange provides a comprehensive solution for supporting resellers. It allows you to give multiple resellers root access to a single shared server, it allows your resellers to manage their customers with a comprehensive customer and billing manager, and it allows your resellers to sell applications such as Microsoft Exchange, Web Hosting, and Miva Merchant.

4. *Ability to offer new services and plans*

Ultimately, you have to provide your customers with all the applications that they demand. An automation solution that is application-independent, and that allows you to host a wider set of applications as time goes by, is preferable to a solution that is limited to a handful of applications. Moreover, a good hosting automation solution should allow you to quickly and painlessly define new hosting plans, and provide the hooks to automatically provision and bill for these plans. ServerXchange satisfies all these requirements.

5. *Ease of use*

With a good hosting automation solution, running a hosting operation becomes much easier. So, you no longer need to pay a large salary to a highly-trained systems administrator. While there will always be a need for a few experienced system

administrators, the large majority of your administrators can be less well-qualified, improving the bottom line. ServerXchange greatly simplifies the difficult tasks of server, application, customer, and domain management. For instance, by making application update as easy and point-and-click, it eliminates a major source of human error, and also makes it unnecessary to hire a highly trained sysadmin to carry out this task.

6. *Integration with existing systems*

Chances are that you have already automated some parts of your hosting operations. A good hosting automation solution should be easily integrated into existing systems. For example, the automation solution may expose the back end using a command line interface. This makes it very easy to integrate the solution with existing systems.

Every action in Ensim ServerXchange that can be managed by a GUI can also be accessed by a command line. This allows ServerXchange to be integrated into your existing systems with little effort.

7. *Scalability*

The hosting market is growing rapidly, and as market penetration is achieved, you will have to deal with issues of scale: how to rapidly serve more customers and more applications on more servers and larger data centers. A hosting automation solution should allow you to scale in all these dimensions, providing adequate support at every step.

ServerXchange allows you to define service plans, and to rapidly instantiate these plans. By opening a reseller channel, it allows you to quickly grow your customer base. In parallel, the comprehensive self-administration interfaces ensures that your support burden does not greatly increase even with a rapid increase in customers. The server and application management features allow you to manage large numbers of servers and applications. In short, ServerXchange allows your business to scale profitably.

8. *OS-independent*

Your customers typically demand services that run on at least two of the three main operating systems in use today: Solaris, Linux, and Windows 2000. ServerXchange supports all three major operating systems.

9. *High density*

You have to deal with a rapidly growing customer base. A high-density automation solution packs more customers into the same rack, thus increasing the return on investment. Lets do some math here. With a standard hosting solution, you can get about 200 sites on a server. If each site pays you about $10/month, you make $2000/month/server. A high density solution lets you host 1000 sites on a slightly more

expensive server. The server cost may go up from $800 to $2500, but you can now make $10,000 a month! You can pay for the more expensive server in a matter of weeks.

With its unique private server technology, Ensim ServerXchange gives you a high density hosting solution today. You can justify its expense using a simple ROI calculation.

10. *Ability to give customers service level agreements*

Since you provide your business customers with business-critical applications, customers can and do ask for guarantees on service availability and responsiveness. A hosting automation solution should help you in setting and meeting these service level agreements.

ServerXchange private servers allow you to give customers and resellers service level agreements. You can modify these agreements using an intuitive GUI.

In short, ServerXchange satisfies all the conditions of a hosting automation solution. In production at over 50 customer sites today, it is the industry-leader. Please visit us at http://www.ensim.com to find out more about ServerXchange, and how you can use it to automate your hosting operations.