Madison River Technologies

Heterogeneous resource and systems management using internet standards

OVERVIEW

The Internet changes everything! The dramatic growth in applications, the proliferation of devices, and the superfluity of services is obvious to anyone with even a peripheral industry interest. What is not so obvious, however, is the complexity and challenges the Internet poses for management of Internet components, applications, devices and services. The coming reality of "connected computing", where everything can be connected via the Internet, creates some interesting and complex challenges for locating, modifying, managing, monitoring, and fault correction of these resources and services. The establishment of the Internet infrastructure, the evolution of several Internet standards, and specific service and agent technology based on these standards enables a new model for powerful and proactive Internet management solutions. In this new model, the Internet is the management infrastructure and the standards are the enabling services. Internet technology enables the following management capabilities:

- Heterogeneous management – Everything (software/hardware servers, workstations, services, infrastructure components, peripherals, remote devices, etc.) managed through common, standards-based methods.
- Auto discovery, configuration and deployment – Internet resources are self-discovered, deployed, auto-configured and updated.
- Proactive management – Policies, directories, and profiles provide intelligence for proactive monitoring, fault detection, and correction.
- Virtual system management – Dissimilar and diverse resources are combined for management in virtual, end-to-end management systems that accommodate interdependencies and relationships.
- Secure remote management – locating, monitoring, and managing all types of resources securely, remotely, and globally.

The benefits to these Internet management capabilities are significant. First, proactive management fundamentally increases reliability and lowers the risks associated with Internet systems. In addition, management is greatly simplified with heterogenous resources being controlled through standard interfaces and methods. The bottom line is lower costs and powerful management. The total cost of ownership, deployment and support is reduced through centralization and standardization of resource management. And the ability to create virtual end-to-end management solutions enhances and extends current management capabilities to include dynamic Internet systems. This white paper outlines the management technologies available and required for the coming reality of "connected computing". It also details the components, services and solutions that Madison River Technologies is providing to enable a new level of Internet resource management. The objective of this paper is to summarize the Internet management problems and detail solutions that will not only solve the current problems but also provide a powerful new class of Internet resource management tools and services.

CONNECTED COMPUTING

Everything Connected

The "connected" thought is not new. Phrases like "the network is the computer", "nothing but Net", and "everything's connected", have been floating through marketing messages for years. What is different is that open standards are now established and any and every vendor can create products and services that are outfitted to work with the Web. The list of Internet-enabled products and services is endless. Just the categories alone include infrastructure components, hardware servers, software servers, applications, client workstations, peripheral devices, remote devices, and many more. Anything that is associated with or connected to a processor is a candidate for networking on the Internet. Here is a brief list--using just the connected world as we know it--of the components that are connected to or work in conjunction with the Internet:

- Infrastructure components – routers, hubs, switches, firewalls
- Hardware servers – motherboards, network adapters, video adapters, storage systems, operating systems
- Software servers – Web, e-mail, news, database, directory, index/search, DNS, DHCP, transaction, e-commerce
- Applications – the commercial list is endless, in addition to custom applications
- Workstations – monitors, mouse, keyboards, CPU+, desktop applications, sound systems, workstation peripheral devices
- Peripheral devices – printers, modems, scanners, storage subsystems, etc.
- Remote devices – laptops, personal digital assistants, monitors, cameras, scanners
- Communication devices – phones, television, settop boxes, projectors, sound systems, cell phones, etc.

Add to this the comprehensive list of other things that have processor power and could be Net-enabled such as lighting systems, automobiles, manufacturing equipment, processing equipment, HVAC systems--the list is endless. What this inexhaustible list points out is the complexity of a networked world.

Each of these resources requires some degree of management; the more standardized the management approach, the simpler and more powerful the management capability. Although many of these resources are currently managed on an individual basis through proprietary methods and interfaces, Internet-standard technology makes it possible to manage resources in virtual systems through common interfaces.

Connectivity to the Internet will soon be as inherent to a device or resource as a microprocessor or memory is today. To use a known analogy, the Internet becomes the bus and all resources are connected components. Through a standard bus, each component is connected and subject to discovery, configuration, updates, and management. Standard protocols provide the common interface for heterogeneous resource management.

Explosive Growth

There is no question about the hyper-growth of the Internet. The following statistics illustrate this rapid, worldwide adoption:

- Annual Internet e-mail growth – 475%
- Annual Internet use growth – 88%
- 10,000 new Internet domain names are added daily
- Total Web-site number is approaching 5 million

In addition to the growth of standard Web technologies, there is rapid adoption and Internet enablement of traditional computer services. Mainframes and legacy applications are being refitted for Internet access. Proprietary applications of all types (Notes, Oracle, Microsoft, accounting, human resources, etc.) are being rewritten to use Internet technology with standard access, standard directory, and standard management capabilities. The evolution of these mainstream information applications to Web standards is also fueling the growth of the Internet.

And finally, there is a complete new generation of Web applications that are being developed from the ground up based on Internet technologies. These applications are accessible through a common browser (HTML, XML, Java), and are distributed, secure, and interoperable with other applications that are also standards-based. These applications, in addition to the new Web devices, peripherals, and services, are fuelling exponential growth of the Internet and Web-related technologies. As with any high growth environment, the biggest problem facing those involved with Internet technologies is management. The need for standardized and centralized management technologies is acute. A recent GartnerGroup quote states, *"Since end-to-end system management does not exist today, opportunities exist for element management solutions that: 1) provide end-to-end management of vendor's product lines, 2) provide detailed information, 3) are relatively simple to implement, and 4) are cost effective."* GartnerGroup's conclusion states, *"Web technologies and CIM (Common Information Model) have sufficient glue to create the technical integration and common services necessary to remove the requirement for a generic framework."*

MADISON RIVER TECHNOLOGIES

Transient Relationships

The Internet also forces a new model of communication and relationships. In the current model, infrastructure is created and relationships are established in a generally well-defined, inflexible and persistent manner. Users can only access a predefined set of printers; access to a database is preceded by manual intervention and account creation by an administrator; network resources must be administered individually. Functions such as self-configuration, load balancing, fault recovery and self-updating do not happen automatically. Internet technologies and protocols allow the formation of ad-hoc and temporary relationships among resources and people. Based on individual characteristics or profiles, resources and people assume different roles depending on different variables such as time, location, activity, and relation to other resources. These roles and relationships are dynamic, ad-hoc and vary according to policies and conditions. For example, a company creates a policy that establishes the price of certain goods (the policy is related to or dependant on other policies or conditions.) A group profile exists with policies that vary depending on conditions. An individual (with their own profile) and with membership in the group will receive personalized, pricing based on conditions at a certain time.

In a simple example, All Sport has an oversupply of soccer balls and would like to selectively unload them to the best potential customers. Since Timmy is online shopping and a member of the Jr. Soccer League, when he hits the All Sport site, he is notified (based on his profile, his group's profile, etc.) that he can get the ball for half price. When the ball inventory hits normal levels, the sale is cancelled and the ad-hoc relationship is terminated.

A second example might include the following scenario. Corporation X has 42 Web servers that are enabled with intelligence about each of the Web server attributes. The company discovers that a security breech is possible if a Web server is running a particular configuration. Using intelligent management, it is possible to query and determine that 31 of the 42 Web servers require a patch to maintain security. Communication is established, and the patch downloaded and deployed on all 31 servers simultaneously.

Another exploding market is personal digital assistants (PDAs) that provide calendaring, scheduling, mail, and task lists, etc. These devices are "stand alone servers" when they're not connected to the network. However, when they connect to a host or server for downloads and updates they become "clients". The transaction or connected state of the device determines the outcome or actions that will flow from the transaction. If the PDA is connected to a central groupware host, the transaction will provide e-mail, appointments, etc. If the connection is to a game host, the transaction may provide new alien worlds for an existing battle game.

PROBLEM AREAS

The problems associated with the Internet can be grouped into several general areas. Technology from

Madison River addresses these problems in a direct, solutions-based manner.

Management

The environment of “everything connected” and “explosive growth” can create a management nightmare. Disparate and dissimilar devices and services usually have disparate and dissimilar management interfaces. The number of different management interfaces companies need just to monitor the components of an e-commerce transaction would include operating systems, Web servers, transaction servers, databases, routers, browsers and all of the hardware associated with each of these services. Each component, service or device requires a separate management interface that requires additional knowledge and expertise. With the growth and distribution of the Internet, management is difficult, to say the least. In addition to different interfaces, all of these resources are managed serially, or one at a time. Each router is modified or managed individually as is each Web server and workstation, etc. There is no accommodation for large numbers of resources that have common attributes or management properties that could be modified in parallel, or all together at the same time.

And finally, there is no ability to manage virtual, systems that are composed of disparate resources but may be grouped together for a specific function or process. The failure of one resource affects the results of the contributions of many other resources. In an e-commerce transaction for example, the success of the transaction is contingent on many different types of resources that are temporarily dependant on each other for the duration of the transaction. The ability to monitor or manage the collective status of hundreds of dissimilar resources is currently not possible.

Scalability

With the global distribution of the Internet, scalability takes on new dimensions – actually in two directions: upward and downward. There is the potential for billions of resources to connect to the Internet. The key to managing large numbers of dissimilar resources is dynamic information. With intelligent information about the configuration, state, policies, etc., of each resource, the network can create directories and virtual management systems on the fly without the need to pre-configure or statically maintain this information. By making resource management information available dynamically, the ability to scale systems increases dramatically. It is no longer necessary to predefine the location, status and authentication required for each Internet resource. Systems are dynamic and automatically accommodate policies that initiate fail over, load balancing, and fault recovery as systems increase in scale.

Downward scalability applies to the myriad of micro-devices and services that are becoming Internet resources. The distributed nature of the Internet makes it possible to granularize and distribute applications and device components. Functionality increases as these components join together as part of a system. Since many components are small and may be lacking or have limited memory, bandwidth, or computing power, intelligent agents and communication make it possible to integrate these micro resources as part of a larger system. It is also possible to segment or reduce large systems into smaller subsystems for easier management. Creating virtual management systems that focus on specific functions, locations, or resource types can enable micro or precision management for low level functionality.

Mobility

In the "everything connected" paradigm, not everything is connected all of the time. Devices and workstations connect and disconnect constantly depending on location and mobility. Personal digital assistants usually only connect for synchronization; laptop computers are disconnected while in transit; home computers only connect while dial-up is established. The "coming and going" nature of devices on the Internet is a reality that contributes to the complexity of management. The management problem that exists is determining a disconnected device's status. Mobile devices routinely experience changes in location, IP address, connected state and connection media. With the current growth of mobile devices and the coming proliferation of special-purpose devices, the problems associated with mobile management will only be exacerbated by the transient Internet connections.

Interoperability

With the diversity of systems, interoperability is always a problem. The Internet, on a basic level, brings together resources of all types but interaction, common management, and functional interoperability becomes the next major obstacle. Some of the problems include getting databases to interoperate with Web servers, management components to interoperate with framework consoles, independent devices to interoperate with standard communication protocols, etc. The Internet has proven that standard protocols are an efficient and effective way for resources to interoperate. Protocols, instead of interfaces, become the interoperability tools. The Internet itself becomes the communication framework and the open-standard protocols and information models are the vehicles for interoperability.

Internet-based Applications

The Internet infrastructure concepts of redundancy, fail over, load balancing, and self configuration are generally lacking when it comes to applications. Applications have generally been developed based on a client/server model where the bulk of application functionality resides at the server and a thin, usually proprietary client sits on the workstation. Using the same Internet principles and technologies (protocols as interfaces, directories, auto-discovery and configuration, etc.) a new class of applications are being developed that have dynamic management characteristics and higher levels of performance and reliability. The Internet itself is the computing model for Internet-based applications. The Internet was designed around the concepts of redundancy, load balancing, scalability and fault tolerance. Building on the e-commerce transaction example, an application built using Internet technologies would have multiple sources for credit card verification; if one source is busy or down, the program would automatically select another based on preferences and policies. A transaction might take any one of several secure routes depending on performance or quality of service level the customer requires. Of course, users could access the application via any standard browser. Specific processes in the transaction such as inventory checking, pricing, order status, etc. could be located and accessed in multiple ways. With modular functionality that is accessible through standard, secure interfaces and protocols, developers can write any distributed application with provisions for redundancy, fail-over, load balancing, and self configuration.

Cost of Ownership

Extensive research from leading analyst organizations reveals the largest component of the cost of ownership of a network is management. GartnerGroup states that 74 percent of the total cost of ownership is administration, technical support and end user support.

Centralization and standardization are the most difficult cost-of-ownership enablers to implement, but will yield the largest potential benefits. Internet technologies have the ability to dramatically reduce these costs. Private studies also show that the cost of deploying a current management framework is actually more than the savings customers gain from the increased manageability. For example, a leading framework vendor derives half its income from deployment services rather than software licensing. Fifty one percent of enterprise companies are dissatisfied with their current costs of network systems management. Ninety four percent consider reducing network management costs important. And finally, there is the often immeasurable but critically important cost of downtime. For online services downtime is a direct cost; for enterprise establishments there is wasted employee time, delayed information flows, and lost or destroyed data.

INTERNET RESOURCE MANAGEMENT SOLUTIONS

The flexibility of Internet technologies enables a vast array of solutions. Combining these technologies can create some powerfully unique management applications. While the list of possibilities is long, Madison River focuses on solutions that solve management problems in the following areas. The following scenarios illustrate the types of solutions companies can create using the technology described in the previous sections.

Discovery, Deployment and Configuration

Problem: Workstation configuration and management are a continual management challenge. System Integrators and corporate IT both have the problem of deploying workstations, keeping them current with software, and supporting them after they are deployed. This expense is the major cost of managing a network.

Solution: Using implementations of Internet technologies and open implementations by Madison River, it is possible to create a centralized and standardized management system that does the following:

- Auto discover a new workstation that has been connected to the network
- Discover the workstations attributes and characteristics (software installed, BIOS, adapter card, etc.)
- Instrument the workstation with agents and software depending on directory profile information relevant to that workstation or control group
- Establish an alerting system to monitor and detect faults
- Based on policies, events and conditions, auto correct faults and auto update resources.

All of this activity happens in a secure environment. Companies can adapt the same technology to discover, update, and configure all types of resources including remote devices, peripherals, applications, etc. Companies can combine dissimilar and diverse resources connected to the Internet for management in virtual, end-to-end management systems. Instead of serially managing the resources from a single vendor or from a single application, companies can manage all resources that impact a particular service or application through a single virtual management system.

Heterogeneous Resource Management

Problem: IT managers administer a growing number of different types of resources. Each of these resources have different management interfaces and must be managed serially. The complexity and large number of different systems, and the differences in proprietary interfaces creates management problems and increases costs.

Solution: Every resource is enabled with Madison River MicroDirectory™Agents or with the CIM [Common Information Model] for identity, status, attributes, state, etc. The management system can auto-discover all resources using SLP and dynamically create a directory of resource information using MicroDirectory™ Services. The information about each resource is exposed through a common management interface such a Web browser. The management system provides operations, management, configuration, etc., through a common user interface where all resources appear consistent and management takes place in the same manner. A hardware server will appear with a menu of configurable options and statistics. The Web server, router, database, etc. will all appear in the same format with similar configurable options and statistics. Consider what is and will be connected to the Internet; the list is endless and ranges from the things that are familiar now, such as servers and clients, to additional resources such as remote devices and appliances. It is critical to manage each of these resources through a common management interface using common standards. It is important and cost effective to gather and manage information about dissimilar resources such as attributes, state, capabilities, dependencies, relationships, etc. in a similar manner.

MADISON RIVER TECHNOLOGIES

Virtual Systems Management

Problem: Internet business processes depend on the end-to-end chain of functioning infrastructure resources. Most resources are managed in silos; each has its own interface and there is no cross-sharing of state information. To manage the infrastructure for a complete business system, it is mandatory that dissimilar devices and resources be combined in virtual end-to-end systems for management with provision for inte dependencies, relationships and state.

Solution: Using standard Internet technologies and information models, it is possible to create virtual system management systems. Companies could combine and expose CIM information about each infrastructure resource through a browser interface. Policy information, maintained in a directory, would dictate what actions happen in response to a particular time or event. The management solution could aggregate and resolve all of the information from a particular system to a few key indicators,.

In the e-commerce transaction example, the following Internet infrastructure resources would be instrumented with CIM data: hardware servers, Web server, transaction server, database server, workstation, routers, directory server, certificate authority, etc.. Each of these resources will have additional attribute and state information that may or may not be critical to a transaction. The "critical" attributes (servers up, bandwidth available, etc.) could be resolved to a few "green lights" indicating that transactions will be successful if executed.

Policies for Proactive Management

The goal of proactive management is to give a system "intelligence". A system should know when to alert, how to react, how to self correct, when to fail over, how to load balance, etc.. Companies establish relationships and interdependencies between management systems and resources using directories, profiles, and policies that predetermine actions and procedures. Internet resources and systems can become proactive and self-managing to increase reliability and simplify management.

Problem: A mission-critical system includes Web, database, directory and certificate servers. Redundant systems are in place but fail over only occurs if there are critical faults.

Solution: Companies can implement a complete virtual system with relationships and actions defined by policies or possibly artificial intelligence. For example, component information from a server exposed through MicroDirectory™ or CIM indicates a CPU temperature rise. This event triggers a level one alert to the console. At level two, the system issues a page or e-mail. At level three, the server fails-over to backup. Companies can interlink systems throughout the Internet with fail-over or load balancing to geographically dispersed servers. Companies can include a resource that is enabled with dynamic state and discovery information as part of a system for intelligent management. Policies moderate the need for micro management and allow intelligent, systems to proactively manage themselves.

MADISON RIVER TECHNOLOGIES

ENABLING TECHNOLOGIES

This new model of heterogeneous resource management is possible due to the evolution and development of several Internet technologies. The Internet itself becomes the infrastructure, the piping, or the framework for management. There is no need to create proprietary, secure, high-speed protocols for management communication; the Internet is already here. Internet standard technologies become the enabling pieces for this new management model. The following table and subsequent detail illustrate the framework and components that exist to create a new management model – the next step is formulating these technologies into management services and applications.

TCP/IP	The common protocol for communication
LDAP	Common directory access for identity and authentication
SLP	Common method for locating Internet resources
HTML/XML	Universal method for access/viewing information
PKI	Common methodology for ensuring security
CIM	Common information model for describing resource characteristics and attributes
DMI	Interface for interrogation of component level resources
WBEM	Web-based enterprise management standards
SNMP	Existing management protocol for status and alerts

Internet Standards

The revolutionary growth of the Internet is due in a large part to the revolutionary approach of standards development. The Internet has defined interfaces in terms of protocols, not in terms of application program interfaces (APIs). What this means is that Internet resources can communicate, interact, and be managed through tools based on common protocols - not vendor specific APIs. Developers create management resource for using protocols, and the protocols (all standard) can work with any Internet resource. A simple example is the management of "Internet appliances". Web servers, e-mail servers, index servers, etc. are all different types of applications. Companies can manage all of these appliances through a common Web browser. There is no need for a proprietary API or interface to each of these services. Information from one may be shared with another - the only common requirement is that a standard protocol such as HTTP be used to access them. The power of Internet protocols as interfaces replaces the need for APIs and opens wide the possibility for common management of dissimilar devices and resources. Internet standard protocols that are critical to management include: Transmission Control Protocol/Internet Protocol (TCP/IP), Lightweight Directory Access Protocol (LDAP) and Service Location Protocol (SLP).

TCP/IP provides the transport and routing functions to ensure that information (data packets) are able to find their way from any Internet resource to any other Internet resource.

Madison River Technologies

LDAP provides a vehicle for accessing and categorizing information about Internet resources. Inherent in LDAP is the ability to define an object's characteristics and attributes and organize them in a schema that other applications and directorires can access and understand.

SLP provides a method for discovering and configuring an object or resource. The combined use of IP, LDAP, and SLP generates some significantly powerful management possibilities. The ability to securely locate, identify, access, and modify any resource anywhere on the Internet - all without the need for any proprietary interfaces - is very appealing.

In addition to Internet protocols, there are other open standards that are valuable in creating Internet resource management capabilities. These include HyperText markup Language (HTML), eXtensible Markup Language (XML), and Private Key Infrastructure (PKI). HTML and XML ensure that access to resources requiring management is global, with standard browsers, and without the need for a proprietary client. Information, no matter what the source or who the vendor, is visible and manageable through a common interface. With XML, the ability to customize and apply formatting and logic to the way this information is viewed is also possible. Public Key Infrastructure (PKI) establishes a secure method for exchanging information within an organization, an industry or a nation. It includes the cryptographic methods, the use of digital certificates and certificate authorities (CAs) and the system for managing the process. Including PKI with the other standards listed here makes the entire management system secure. Resources and attributes are only visible or manageable by authorized applications. Access to a particular resource is controlled. Modifications, configuration, and deployment is possible only upon proper verification.

Information Models

Also part of the standards effort is the establishment of information models for unified categorization, organization, and access to information about a particular device or resource. Current information models that enable management include: Common Information Model (CIM), Web-Based Enterprise Management (WBEM), Desktop Management Interface (DMI), and Simple Network Management Protocol Management Information Bases (SNMP MIBs).

CIM is implementation independent, allowing different management applications to collect the required data from a variety of sources. CIM includes schemas for systems, networks, applications and devices, and new schemas definitions are possible. It also provides mapping techniques for interchange of CIM data with MIB data from SNMP agents and MIF data from DMI-compliant systems.

WBEM is an umbrella term for using Internet technologies to manage systems and networks throughout the enterprise. Companies can use browsers or applications to access the information that is available in formats such as HTML and XML.

DMI provides a bi-directional path to interrogate all the hardware and software components within a PC. When PCs are DMI-enabled, their hardware and software configurations can be monitored for status and

SNMP agents are hardware and/or software processes reporting activity in network devices (hubs, routers, bridges, etc.) to consoles that oversee the network. The agents return information contained in a MIB, which is a data structure that defines what is obtainable from the device and what can be controlled.

Using a combination of Internet protocols, these existing information models and the resource information associated with them it is possibleto create comprehensive and powerful management tools. Once a resource is identified, a management system can expose all of its characteristics and properties through standard information models for management and modification.

Symmetry

The interconnectedness and intelligence of Internet technologies provides unique identity characteristics. With Internet technologies, the concept of symmetry or dynamic role changing is possible. At the simplest level, a child process can become a parent or a parent can become a child, depending on events or policies. For example, a Web server can publish data (server) and also request data (client), thus being a client or server depending on the situation. A resource can assume one of several roles depending on events or policies. Symmetry is a fundamental but little understood capability of the Internet. Originally designed to withstand failure to any part of the system, the Internet has the inherent ability to repair and regenerate itself. A good example of this concept is the Domain Name System (DNS). With interconnected DNS servers located throughout the Internet, if one supplying server goes down, a receiving server can become the supplying server and can avert a failure of the system. The fact that services are interconnected, can respond to events, and change role based on conditions is at the heart of the Internet,s robust nature and at the same time provides some powerful management capabilities. An example of this is a process that requires parent, child, and peer services from a multi-tier application. Any child service can become a parent or peer service depending on the conditions that exist. A service may be requesting data, providing data, or manipulating data depending on a policy or condition. Managing or stipulating these changes manually is far too tedious. Furthermore, Internet resources do not remain in a constant state, but are constantly changing. This concept becomes very powerful when exploring critical management issues such as load balancing, fail-over, redundancy, fault detection and correction. The combination of open communication protocols with events, policies and profiles makes it possible to create easily managed systems that are robust, dynamic, and intelligent. Discovery, Deployment and Configuration Services. With an infinite universe of possible Internet services and resources, just locating and organizing them is a big problem. Once a resource has been found and added to a "system" it is often necessary to configure it to work within the system or deploy updated code. The ability to automatically find, automatically update, and automatically configure a resource is extremely valuable with the current acceleration of Internet resource additions.

To solve this problem, the Internet Engineering Task Force (IETF) has ratified Service Location Protocol version 2 (SLP v.2). SLP makes it possible to automatically find resources (hardware, software, services, etc.) anywhere on the Internet. A device or resource can self register with a device or service registry. Again, using profile and directory information, a device can determine its proper configuration and update requirements. This provides makes it possible to dynamically configure a resource based on events or conditions. A management system can store configuration information for a device in multiple locations, so that if one location is inaccessible, the system can retrieve the information it needs from another. With the ability to discover, establish two-way communication, and initiate operations with various resources comes powerful management capabilities. The most obvious is software distribution or resource instrumentation. Companies can automatically configure workstations with applications or to specific preferences. Companies can update all types of resources with current drivers or patches.

A good example of this is a digital camera with cell phone capabilities. When a user turns the camera on, it establishes an IP connection through the phone line. On connection, a camera agent alerts its presence using SLP. Information about the camera such as current BIOS, drivers, and configuration statistics populates a directory service. Based on an existing policy (which states that the drivers prior to a certain date need replacing) the management system deploys new drivers to the camera. Management of the camera such as image interval, phone number to dial, etc. is handled either manually through a browser or automatically through profiles that are stored in a directory. From time to time, it may be necessary to update the phone software. The Internet connection can handle this automatically. This can all happen while the camera and the hosts are at different locations anywhere in the world. Since the camera is only actually part of the network while it is connected, the fact that it can self register and auto-update eliminates the need for manual monitoring or status checking. The result is a dynamic directory of resources with no need for static databases that can't keep up with change, growth, or connected status. The device knows the information about its current state and attributes and updates the directory transparently and dynamically when it connects.

Directory Services

Directory services, though they have been around for years, are still one of the least understood yet most powerful mechanisms for management - especially on the Internet. Fundamentally, a directory service does two things for an Internet resource: provide identity and ensure security.

The "identity" of a resource can be a complex thing and a true directory service is designed to address this. Among other things, name, location, relationship to other resources, common characteristics, unique characteristics, and even source code can be subsets of directory service information. A directory service also enables security. Relationships of one resource to another, resources to administrators or users, and resources to other directories can be established and security enabled based on these relationships. Companies can accomplish the management task of who and which groups have access to what with a directory service. Determining identity, location, and access is critical for any resource or service available on the Internet. With the evolution and development of directory services over the Internet, management systems can identify any resource - whether it be a hardware or software server, a router, a peripheral, an application, etc. - and can quickly identify what it is and what its characteristics are. Directory services also make it easy to locate the resource and securely authenticate it or allow access.

The most powerful aspect of a directory service is the ability to manage complete groups or classes of attributes with a single operation. If you have 300 workstations, each with an application that requires a Y2K patch, initiating the process to install one patch is all that is required to install 300 patches. The directory knows the "patch fix" state of each machine and can manage all of them because they have a single common attribute. The system can execute operations in parallel rather than serially. The main disadvantages of all current directory services are that they are static in nature, not easily extensible, and must be manually generated. There is no ability to automatically determine and correct for resource connection or state. It is arduous to add new classes and types of resources. Worst of all, there must be some sort of manual intervention to include a new resource as part of the directory.

Security Services

Security services on the Internet have been a major source of confusion and pain. Companies are inherently skeptical of Internet security. ThePublic Key Infrastructure (PKI) provides a method for ensuring that access to a particular resource is only allowed by an authorized agent.

Although PKI has been standardized and is in use, there are limitations that restrict widespread use and adoption. These limitations include the general availability and distribution of public keys that would make them commonly available to every resource that could take advantage of them. There is also the problem of maintaining lists of certificate authorities and current certificate revocation lists.

However, the combination of PKI with directory services and SLP eliminates these barriers. With slight modification, it is now possible to piggyback distribution, revocation, and maintenance of PKI on discovery and directory protocols. This technology allows companies to implement PKI at the level that it was originally conceived and provides a comprehensive security system for all resources on the Internet.

MicroDirectory™ Technology

Combining the above developments with MicroDirectory™ technology from Madison River enables a new and powerful class of management options. The MicroDirectory™ provides dynamic and intelligent options for management, fault prevention, detection, and correction. At a basic level, the MicroDirectory™ intelligently enables any Internet resource for management. Once enabled, a resource can be auto-discovered, controlled through policies, configured as part of a logical system, integrated into other management frameworks, or managed through one of several methods. Resource information that is available through standard information models is utilized and augmented. Simply put, the MicroDirectory™ enables a resource with intelligence for management.

INTERNET RESOURCE MANAGEMENT BENEFITS

The benefits to the forms of Internet resource management described in this paper are immense. Going forward they become more measurable and valuable as "connected computing" becomes more of a reality.

Generally, the value propositions to Internet-based resource management can be summarized in increased reliability, lower risk, simplified management, higher availability, greater scalability, and lower cost of ownership and management.

Increased Reliability

Up-time and freedom from critical failures are direct benefits received from solutions built using Internet resource management technology. Self monitoring, self correcting, auto load balancing, and auto fail-over capabilities all work together to eliminate catastrophic or mission critical failures. This translates into lower risk for any system that is part of an integrated Internet network.

Simplified Management

The use of a common browser interface, common information models, and standard methods for directory and policy management greatly simplify the task of management. Using Internet resource management technologies makes it possible to unify disparate and heterogeneous systems for common management – they all are different, yet their management is the same. It is also possible to use directories to leverage the amount of management effort required. Companies can manage an entire group or class of resources with the same number of steps as a single element. Diverse and heterogeneous resources are aggregated to a common point for seamless, consistent management.

Madison River Technologies

Enable Internet Scalability and Availability

A major benefit of Internet resource management technologies is increased scalability. The Internet is an unfathomably large network. However, the technologies described in this paper make it possible to distribute services, applications, and management to the extents of the Internet. Scalability, when addressed with proper location, distribution and management technologies, is not an issue. It is the same with availability; fail-over and load balancing when properly implemented and managed increase the level of availability of Internet resources.

Keeping Consistent with the Internet

In "Internet time", events happen exponentially faster. Given proliferation and growth, it is imperative to have an existing system "stay current". Auto discovery, deployment, auto configuration, and auto update are all functions possible with Internet resource management technology. Keeping up on the number of new browser plug-ins and Internet software updates, ensuring that applications have the latest new features, updating firmware, etc. are all processes that companies can automate to keep an organization consistent with the latest technologies.

Lower Cost of Ownership and Support

Much has been written about lowering the cost of ownership for network technology, However, the costs to administer and maintain a network still remain high. Internet resource management technology significantly lowers cost in several ways. First, the availability of common, low cost, simple function, easy to use, open Internet technologies provides the infrastructure for new low cost management. Existing framework system management solutions require the implementation of their own infrastructure; the Internet becomes the infrastructure for Internet resource management and thus eliminates the high implementation costs. The pieces companies and solutions providers need to create an Internet-based management solution are generally low cost or free. Indirect cost savings result from all of the benefits this paper addresses such as increased reliability and scalability, simplified management, lower risk and less down time. Regardless of the area of focus or implementation, Internet management technology can be a catalyst for reducing costs by resolving variables to a common level and providing easy access, unified interfaces, and consistent methods for management.

MADISON RIVER TECHNOLOGIES

SUMMARY

In conclusion, a new management model is emerging - quickly. It is being propelled by the same fundamental principles that are responsible for the current high growth in the Internet. This includes the ubiquity of common, open standards that are granular enough to provide mass adoption and a wide spectrum of uses. Madison River Technologies is focused on catalyzing the adoption of Internet standards and creating solutions that significantly enhance Internet resource management. The building blocks to these solutions include SLP, LDAP, WBEM, XML, CIM and MicroDirectory™ technology. Management solutions based on these technologies will globally and securely provide the following:

- discovery, deployment, configuration and updating
- heterogeneous system management
- virtual system management
- proactive policy management

The result will be more reliable Internet systems that are simple and less expensive to manage. These solutions ensure end users availability, performance, and more access.