

Orinda, California

High Availability Enterprise NT Infrastructure Design and Implementation

Deploying Microsoft Cluster Server for the Enterprise

Background

In a recent major business initiative, one of the nation's largest healthcare providers faced a daunting challenge: how to overhaul its aging computer infrastructure to handle the complexities of a national enterprise serving 9,000,000 members. With almost 70,000 workstations distributed across the United States, the company struggled with the same difficulties of other organizations its size: how to distribute and support – in a timely, cost-effective manner – the myriad of applications required by its users. Moreover, the company also needed to ensure continued system availability and update its failover processes to ensure high resource availability for mission-critical applications across the enterprise.

The infrastructure upgrade contemplated by the company constituted a massive undertaking involving the design and construction of a distributed computing infrastructure based on Microsoft's NT architecture, but enabling co-existence with UNIX and NetWare platforms. Recognizing that it needed to augment its internal resources to carry out a project of this scope, the company engaged Axean Pacific, a technology consulting services firm based in the San Francisco Bay Area, to design and architect an optimized, high-availability, scalable solution.

A Multitude of Challenges

Axean Pacific's overarching challenge was how to design, deploy and maintain an NT desktop and server infrastructure that supported core enterprise applications and local market initiatives, while ensuring high resource availability and manageability of client data, and lowering overall costs. But a series of other challenges quickly came to the forefront. For example:

- The project's scope included providing centralized support for decentralized deployment, and involved large-scale simultaneous installations and migrations from the legacy to the new environment.
- A "managed" desktop – capable of providing access to enterprise applications and file and print services – was a base requirement of the project. These core service requirements expanded further still to include asset management, software distribution, name resolution, address management, quota management, internet services, NetWare and terminal emulation gateway services (for mainframe and Unix hosts), and backup and recovery.
- In order to provide all the services outlined above, the infrastructure design would require a core set of servers. To optimize hardware usage, utilize network bandwidth efficiently, and ensure high-availability, the core servers would have to complement each other's roles within the infrastructure. Timing was a critical factor as well: The organization's fiscal year budget constraints dictated that a server solution, once finalized, would have to be rolled out to over 150 sites within a 90-day time period.
- Cost was a primary driver as well. Axean Pacific's design, once implemented, would have to result in a decrease in overall computing costs and improved system manageability and availability. Availability was measured not only in the users' desktop usage in their primary location, but across locations that users were expected to visit in a given workweek. The losses that resulted from unavailable systems far outweighed any other cost factor. These objectives, in turn, would have to be balanced with other business requirements, such as autonomy at the enterprise's various sites and centralized control.

The Solution

After completing the requirements analysis and project planning, Axean Pacific developed several “proof of concept” scenarios based on an engineered approach to meet the company’s requirements. It was given that the technical effort required to successfully execute the project would be significant, but its organizational aspects would demand considerable management effort as well. Therefore, validation of the proof of concept scenarios would have to include verifying the management aspects of the project, as well as the effectiveness of the system design from the standpoint of implementation and support.

The infrastructure design that was ultimately accepted by the company proved successful, and the company currently employs Axean Pacific to manage the transition of its ongoing systems support and the expansion of its infrastructure into other areas of the enterprise.

The design elements include:

- High availability systems, from the desktop to back office support systems
- Standard desktops across the infrastructure for ease of software distribution and remote support
- Scalable infrastructure and deployment methods
- A support model that provides centralized oversight, but allows for local market execution
- An approach for leveraging hardware vendor relationships that lowers costs while increasing availability and geographic coverage
- Effective configuration management and change control
- A sustainable engineering effort that results in a predictable and supportable environment
- Proactive system monitoring and management

The above design incorporates within it (1) swift procedural and automated deployments, (2) reduced support costs, and (3) improved availability – which collectively resulted in our reaching the earlier stated goal of reducing overall operating costs.

The Approach

To carry out this project, Axean Pacific employed a phased methodology designed for developing and implementing enterprise-level technology infrastructure upgrades. Every aspect of the project – from assessments of the current architecture to proof of concept scenarios, development of the engineered solutions, implementation, and ongoing support and review – followed the tenets set forth in the methodology.

A first step in the design process involved identifying common components – such as high-availability functionality from the desktop to the server -- and developing them for re-use throughout the project. In addition, core teams (organized by geographic area) were formed to handle the project’s complex organizational aspects – from design to technical training, information management, installation and deployment. Team members employed the methodology in developing and validating their work product and ensuring close coordination with other teams.

Three additional project elements also contributed to its success:

- Decision-making bodies (called “governance boards”) were set up at multiple levels of leadership within the enterprise. The governance boards assisted with overall project management, including local market acceptance of the project.
- Centralized project controls for finance and planning were implemented to assist with the management of project costs and deadlines.
- Detailed “Statements of Work” – a standard documentation deliverable – were developed to clarify activities, costs, staffing, and scheduling associated with each aspect of the project, from engineering to deployment and support.

Technology Solutions

Various factors contributed to the decision to utilize specific products for the project's technology solution. Technical and functional merits were balanced with cost and installed base factors. The technology solution that we arrived at included the following hardware and software components for the desktops and servers:

Figure 1: Desktop Solutions

APPLICATIONS	CATEGORY
Microsoft Office	Office Tools
Paradox	Database/Transitional App
WordPerfect	Office Automation/Transitional App
Harvard Graphics	Presentation/ Transitional App
Visio	Graphics
Lotus Notes	Electronic Messaging
Adobe Acrobat	Document Viewer
Microsoft Internet Explorer	Web Browser
QuickTime for Windows	Multimedia
Attachmate	Host Access
McAfee VirusScan	Anti-Virus
Microsoft Project	Project Management
Crystal Reports	Report Generation

Figure 2: Server Solutions

APPLICATION	CATEGORY
NTS Enterprise Edition	Operating System
MS Cluster Server	Storage Management/Failover
NT Resource Kit	OS Tool
SQL Server	Database
MS Internet Information Server	Web Server
QuotaManager	User and Group Storage Administration
Seagate Backup Exec	Storage Management
Enterprise Administrator	Security and Administration
Intel LanDesk Virus Protect for NT Server	Anti-virus
SNMP	Systems Management and Support
Compaq Insight Manager	Systems Management and Support
HP OpenView	Systems Management and Support
Systems Management Server	Software Distribution
MS Gateway Services for NetWare	NetWare Connectivity Tool
MS File and Print Services for NetWare	NetWare Connectivity Tool

Engineering, Application Integration and Deployment

Following one of the basic tenets of Axean Pacific's methodology, all of the project's engineering deliverables were designed and modeled prior to being moved into production. Working within the lab environment, our engineers developed automated build routines for each server type as well as the standard desktop. The automated routines were designed to meet one of the most critical objectives of the project: the ability to support multiple, simultaneous deployments at more than 150 sites across a wide geographical region. To better leverage the use of vendors in the installation of hardware, the build routines had to be designed to run unattended for most, if not all, of their phases.

Working with the engineers, Axean Pacific's technical writing staff developed detailed checklists ("job aids"), procedures, and planning guides designed to reduce the need for highly-skilled engineers during the deployment phase. In addition, Axean Pacific provided the organization's IT staff and vendor technicians with comprehensive training on all aspects of the deployment process – from site preparation to server/desktop deployment methods to final quality assurance and production acceptance.

As with the server build engineering, applications were integrated centrally and deployed according to local area needs and schedules. A combination of client- and server-based applications were deployed to support the organization's requirement for a roaming desktop. Support was likewise defined across sites using a tiered model for escalation and elevation of issues. A web-based Support Desk tool integrated with Systems Management Server (SMS) was used to track the project's progress as well as individual service requests and support calls.

MS Cluster Server provided a 2-node failover system that enabled 24x7 availability for file and print services. (For a complete overview of MSCS, see *Providing a High-Availability Solution* section below.) Since user files and profiles resided on the file servers (in order to support full roaming and file recovery throughout the enterprise), the introduction of clustering technology was crucial to the success of the project. In addition, server roles were optimized across the enterprise to provide primary and multiple secondary resources for all core services -- including authentication, WINS and DNS name resolution, and asset management and software distribution using System Management Server (SMS). DHCP was implemented using split scopes on separate LAN segments and servers for each site.

Utilizing proprietary tools developed by Axean Pacific, our engineers developed centralized login scripts and implemented policies to manage the "look and feel" of the desktop, as well as control access to applications. Our engineers also created routines that, in the event of multiple system failures, would automatically locate the nearest application server and a backup machine. Core file and print services included (1) transition support from NetWare servers, (2) interoperability with UNIX, and (3) gateway services to the mainframe. In order to enable full roaming for notebook PCs, EOTK routines were implemented to support synchronization of notebook files.

Providing a High-Availability Solution: MS Cluster Server

As mentioned at the beginning of this article, high resource availability was identified as a critical element for this company's infrastructure upgrade project. While down times of even a short duration are unacceptable in most large companies today, a system failure within a medical facility may constitute even greater consequence. In part due to its significant expansion in recent years, the healthcare organization's initial failover solutions no longer provided the up time and recovery required by its mission-critical applications. The release of Microsoft's Cluster Server (MSCS) product provided the answer to the organization's problems.

Like other clustering technologies before it, Microsoft Cluster Server enables an enterprise to connect a group of servers in order to improve data availability, server manageability, and performance. Even with two physical servers connected in a cluster, the workstation responds to the cluster as if it were a single server. In the event of a system failure, cluster software automatically disperses current workloads from the failed system to remaining systems in the cluster, thus restoring user access to data and services without interruption. In fact, end users are typically unaware that a failure has occurred. Hoping to improve upon clustering solutions of the past – which were often complex, difficult to configure, and relied upon expensive proprietary hardware – Microsoft is developing MSCS for its NT server-based OS based on open specifications and industry-standard hardware.

Despite certain limitations (see *Tips and Techniques* below), MS Cluster Server proved to be an excellent solution for providing core desktop services for the healthcare organization's vast network of workstations. Prior to deployment, Axean Pacific's engineering team put the product through its paces using a variety of scenarios and it performed solidly with respect to failover and failback. All in all, the product provides an excellent 2-node failover solution for NT enterprise infrastructures that require high-availability systems.

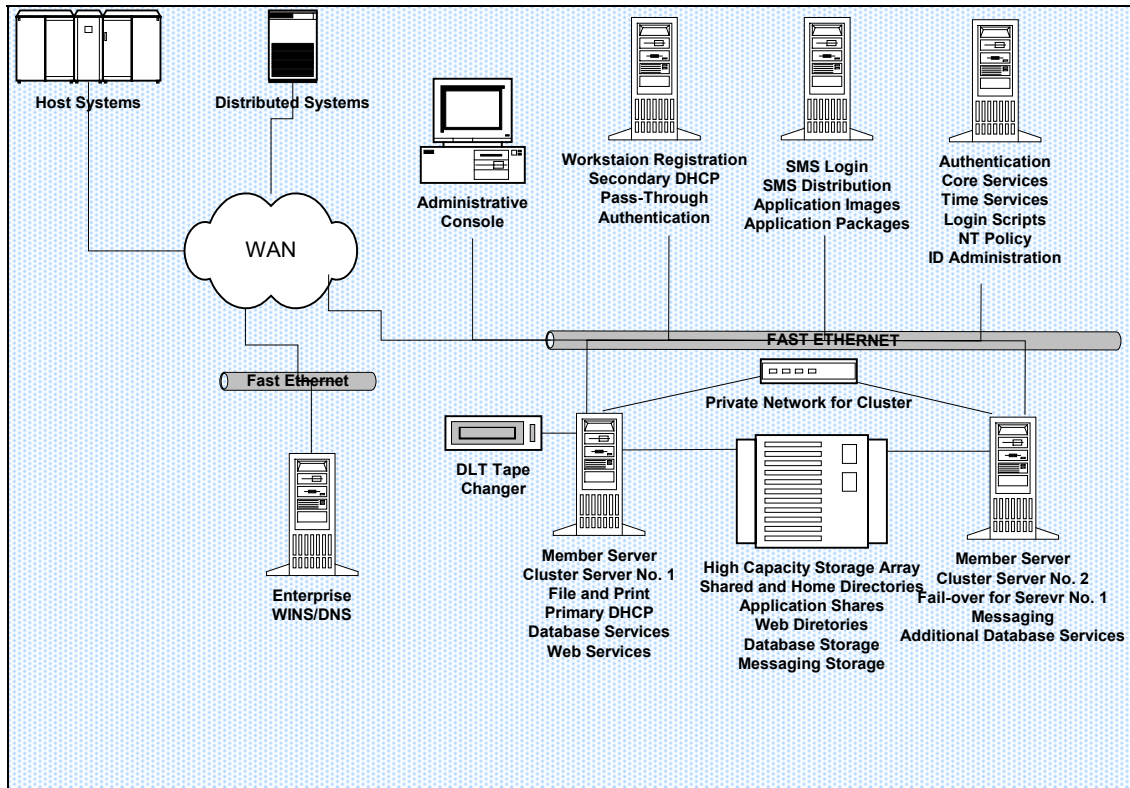


Figure 3: High Availability NT Infrastructure

Deploying MS Cluster Server: Tips and Techniques

While MS Cluster Server represents an improvement in the high-availability arena, Microsoft's claims that MSCS is easier to configure and implement than older clustering technologies are not completely realized in Release 1.0 of the product. To ensure the success of the teams responsible for installing MSCS at multiple sites, Axean Pacific developed detailed procedures for dealing with the product's relatively complex hardware dependencies and software set up processes. In addition, during the project's Design & Modeling Phase, we found it necessary to develop several tools specifically designed to address some of the current limitations of MSCS.

For example, our engineers developed a tool to simplify the rather complex process of configuring administrator file and print shares. Another tool tied the creation of home shares to user identification and group administration, and then integrated the entire process into Enterprise Administrator. We engineered other tools simply to expand upon the solutions provided by MSCS. Examples of supplemental tools include one that migrates existing shares (on existing servers) to the cluster configuration, another that automates the retention of file permissions, and a third tool that recreates existing print queues and definitions.

In addition to the limitations outlined above, Axean Pacific also noted several areas that will hopefully be improved upon in future releases of the product. For example, the v1.0 release of MSCS provides relatively limited support for storage hardware. In addition, the applications that support MSCS are fairly limited, although it provides support for core services like IIS, file and print, and DHCP. (Ironically, MS SQL Server v6.5 was not cluster-enabled at the time of the release of MS Cluster Server v1.0.) In our

tests, Oracle v7 and v8 worked well with MSCS, as did LotusNotes. It should be pointed out, though, that proper testing must be conducted prior to deploying MS Cluster Server to ensure that all applications work well within the cluster environment. Another significant consideration involves the re-addressing of physical and virtual servers in a cluster. Changing IP addresses is not a straightforward process and currently requires reinstalling MSCS, a prospect that is untenable for existing installations. Finally, Axean Pacific noted difficulties with the scalability of MSCS and intermittent problems with permissions – again, issues that we trust will be addressed in future releases of the product.

Supplemental Tools Developed by Axean Pacific addressed limitations with Microsoft Cluster Server in the following areas:

File and Print Share Configuration ~ Home Share and UserID Configuration ~ Migration of Existing Shares to the Cluster Configuration ~ Automatic Retention of File Permissions ~ Automatic Re-creation of Existing Print Queues and Definitions

High Availability for Windows 2000

There are a number of alternatives to support high-availability for application servers. These alternatives include hardware-based solutions for redundancy, hot-standby, from hardware components to whole server setups. Application-based solutions include load balancing to fail-over solutions for specific applications or application components. In Windows 2000, there are a number of new features, as well as enhancements from those available in Windows NT 4.0 and MS Cluster Server v1.0. The following summary discussion presents these options and provides a background on how the solution discussed in this white paper can be upgraded using Windows 2000.

Windows 2000 high-availability solutions will ultimately be compared to UNIX-based solutions and, with Windows 2000, the Microsoft offerings become very competitive in price and features. With Windows 2000, scalability and high-availability features now play in the large-scale memory space.

Advanced Server supports up to 8GB of memory and SMP support for up to 8 processors. Windows 2000 Datacenter, set to be released later in 2000, will support up to 64MB of memory and up to 32 processors. Windows 2000 Datacenter, when released, will not only equal but surpass most UNIX-based solutions in terms of large memory and SMP support.

Essentially, Windows 2000 offers Load Balancing and hardware-based fail-over using Cluster Services. Load Balancing is supported in Advanced Server by distributing requests for IP-based application services across participating servers. For example, web server requests are managed through an application server running Windows 2000 with Network Load Balancing. This server, in turn, identifies available web servers, and appropriately distributes requests across the participating servers. While supported across 2 servers in Windows NT running Cluster Services, Advanced Server now supports this feature across 32 participating web servers in a cluster.

For hardware fail-over, MS Cluster Services still only supports two-server fail-over but with enhanced management features. Four-server fail-over will be supported in Windows 2000 Datacenter. One thing still remains - be sure to check the MS hardware compatibility list for supported hardware. This is crucial since not all platforms for both servers and storage is supported. As with Windows NTS and any other high-availability solution, the combination must be thoroughly tested prior to deployment. Application component-based dynamic load balancing will be supported when MS ships AppCenter - which is expected to provide support for load balancing one to many COM+ components for servers participating in a cluster. With this offering, MS will have offered solutions for all tiers of application servers.

Conclusion

Most of today's enterprise-level infrastructure projects require rapid deployment to accommodate the introduction of differing levels of service and tools, and other project management challenges associated with a phased deployment approach. This project was no different. With core engineering nearing completion, the project's deployment phase was defined in approximately one month. This phase included details on scope, locations, staffing, hardware and software sourcing, vendor selection and training, and internal coordination.

Axean Pacific, utilizing a methodology specifically designed for IT infrastructure projects, effectively leveraged technology products to meet real-life business needs. The combination of technical project management, scalable processes and automated routines built from our toolkit enabled Axean Pacific to respond to this project's requirements. Our success can be measured by the project's results: the actual deployment was executed in just over two months and involved rolling out 432 servers in 112 locations. Of those 400-plus servers, 160 participate in a cluster – which, we believe, constitutes the largest MS Cluster installation to date in the United States.

If you are interested in learning more Axean Pacific and its services, contact Victor Tayao (victor.tayao@axeanpacific.com) or visit our website at www.axeanpacific.com.