



International Business Machines Corporation

*Street Address
City, State ZIP*

Date

Mr. Customer Name

Title

Customer

Street

City, State Zip

Dear Mr./Ms. Customer:

Thank you for the opportunity to present IBM® TotalStorage® SAN File System (SFS) for your file system and data management needs.

Storage area networks (SANs) enable you to share homogeneous storage resources across the enterprise. But like many companies, your information resources may be spread over a variety of storage and server environments, often using products from different vendors. This can result in a multitude of file systems that need to be managed individually, which can increase complexity, require a variety of skilled resources, and make it difficult to implement consistent policies for file and database management. Also, file and data administration tasks often impact application availability. All this can lead to poor utilization of storage resources, high costs, and reduced business efficiency. Why? Because many of today's file systems do not exploit the benefits of SANs.

The SAN File System is designed to help lower the cost of storage management and enhance productivity by providing centralized storage and data management for supported servers using the SAN File System. It provides policy-based storage management through the use of customer-defined rules that specify what storage is used when a file is created. With SAN File System V2.2, you can define rules that specify when a file should be moved between different storage pools and also can specify when a file should be deleted. Storage can be virtualized into storage pools organized by service class, and files can be automatically placed on the appropriate storage resource, thereby, helping reduce administrative overhead.

The attached proposal reflects our desire to help ****CUSTOMER**** deal with file system and storage management challenges by leveraging the benefits of IBM TotalStorage SAN File System solution. This proposal demonstrates our long-term commitment to a successful relationship with you in striving to achieve your business objectives. We are ready to help you increase your storage return on investment and reduce your costs and risk of application downtime.

Sincerely,

Name

Title

Telephone, e-mail address



Customer Name

**IBM TotalStorage[®]
SAN File System
Proposal**

***Contact Name
IBM Corporation
Address
Phone Number***

Date of Proposal

Table of Contents

EXECUTIVE OVERVIEW.....	1
WHY IBM TIVOLI?	2
IBM PROFESSIONAL SERVICES	2
IBM TOTALSTORAGE SAN FILE SYSTEM	3
WHAT IS IBM TOTALSTORAGE SAN FILE SYSTEM?	3
MAIN FEATURES	4
SAN FILE SYSTEM ARCHITECTURE	7
EXAMPLE OF IBM SAN FILE SYSTEM EXPLOITATION - DATA BACKUP AND RECOVERY	9
IBM SAN FILE SYSTEM COMPONENTS	10
<i>IBM SAN File System Administrative Interface</i>	11
IBM SAN FILE SYSTEM SECURITY.....	11
SERVICES FOR IBM TOTALSTORAGE SAN FILE SYSTEM	13
STG IMPLEMENTATION SERVICES FOR SAN FILE SYSTEM.....	13
STG MIGRATION SERVICES FOR SAN FILE SYSTEM (NON-DISRUPTIVE DATA MIGRATION).....	13
PRICING SUMMARY.....	15

Figures

Figure 1 - SAN File System architecture	8
Figure 2 - Exploitation of SAN File System with ITSM	10

Tables

Table 1. SAN File System user roles 12

Executive Overview

:Business Partner: Insert your own introductory text here and edit text below if necessary. Be sure to include content on your added value here and/or in a later section.

Thank you for giving us the opportunity to share our IBM® TotalStorage® SAN File System solutions with **CUSTOMER**. We are very interested in becoming your partner in managing storage, storage utilization, and storage costs. We feel that SAN File System meets **CUSTOMER**'s requirements. In this proposal, we will demonstrate how SAN File System can help optimize and facilitate the management of your enterprise storage environment.

Most important to your organization would be that the SAN File System (SFS) solutions portfolio could help you increase your storage ROI by enabling you to:

- Improve Storage Utilization
- Reduce Storage Costs
- Manage More Storage with the Same Staff
- Increase Application Availability

SFS is designed to simplify file and data management in your storage area network by consolidating file systems across UNIX®, Microsoft® Windows®, and Linux® servers. With SAN File System, files and file systems are viewed and managed as a centralized IT resource with a single point of administrative control. By providing a common SAN-wide file system, SFS can support secure, heterogeneous data sharing and centralized policy based storage management in an open environment.

SAN File System is designed to help you achieve data lifecycle management efficiencies through policy-driven automation and tiered storage management. User-defined policies provide the ability to better match the cost of your storage resources to the value of your data. SFS is also designed as a highly scalable solution supporting both very large files and very large numbers of files without the limitations normally associated with Network File System (NFS) or Common Internet File System (CIFS) implementations. These capabilities contribute to reduced storage management costs, enhanced productivity, and optimized storage resource utilization.

The IBM TotalStorage SAN File System is designed on industry standards so it can:

- Allow data sharing and collaboration across servers over the SAN with high performance and full file locking support, using a single global namespace for the data.
- Help provide more effective storage utilization by reducing the amount of duplicate data and by sharing free and temporary space across servers.
- Help improve productivity and reduce the “pain” for IT storage and server management staff by centralizing and simplifying management through policy-

based storage management automation, thus significantly lowering the cost of storage management.

- Facilitate application server and storage consolidation across the enterprise to scale the infrastructure for storage and data on demand.
- Help simplify and lower the cost of data backups through built-in, file-based FlashCopy[®] image function.
- Virtually eliminate data migration during application server consolidation, and help reduce application downtime and failover costs.

Why IBM Tivoli?

Over the past decade, IBM has been awarded more patents than any other company. According to data compiled by MicroPatent LLC, IBM received more than 15,000 patents during the 1990s, triple the output in the 1980s total and 2,300 more than the second-place runner-up. IBM's broad patent portfolio includes all of the Tivoli[®] Storage solutions, and is complementary across all of IBM's server offerings. IBM is a leader in the development and implementation of total cost of ownership (TCO) related technologies. IBM has the understanding and implementation expertise of "best practices" that, when deployed with the technology in integrated fashion, can yield an efficient, business-focused IT infrastructure. IBM, along with Tivoli, has the ability to leverage the best in technology and contribute to reduced TCO. The result is an unbeatable combination with solutions in storage that are designed with reliability, availability, scalability, and manageability in mind.

Together, with Tivoli and IBM solutions, we are ready to help you reduce your system costs and risk and to increase your storage return on investment today. We look forward to the opportunity to work with ****CUSTOMER**** to demonstrate these potential benefits.

IBM Professional Services

IBM Software Services for Tivoli provides worldwide technical consulting services for Tivoli Software products that help make it easy to design, build, test and deploy a solution to manage your on demand business infrastructure.

IBM is dedicated to the success of your technology management solution. With a focus on technical consulting for the Tivoli software products, IBM Software Services can provide you with deep technical product skills, solution architecture, knowledge transfer, and quality assurance.

IBM TotalStorage SAN File System

What is IBM TotalStorage SAN File System?

IBM® TotalStorage® SAN File System (SFS) is a SAN-based distributed file system and storage management solution that enables many of the promises of SANs, including shared heterogeneous file access, centralized management, and enterprise-wide scalability. In addition, SFS leverages the policy-based storage and data management concepts found in mainframe computers and makes them available in the open systems environment. IBM TotalStorage SFS can provide an effective solution for customers with a small number of computers and small amounts of data, and it can scale up to support customers with thousands of computers and petabytes of data.

Customize the following paragraph based on your customer situation:

... can provide an effective solution for ****CUSTOMER****'s "X-Application" consisting of X host servers and supporting X TB's of data. As ****CUSTOMER****'s needs grow and expand with additional application and data requirements, your SAN File System can scale to support ...

The SAN File System has been designed as a network-based heterogeneous file system for file aggregation and data sharing in an open environment. As a network-based heterogeneous file system, it provides:

- High performance data sharing for heterogeneous servers accessing SAN-attached storage in an open environment.
- A common file system for UNIX® and Microsoft® Windows® servers with a single global namespace to facilitate data sharing across servers.
- A highly scalable out-of-band solution supporting both very large files and very large numbers of files without the limitations normally associated with NFS or CIFS implementations.

IBM TotalStorage SAN File System is a leading edge solution that is designed to:

- Lower the cost of storage management
- Enhance productivity by providing centralized and simplified management through policy-based storage management automation
- Improve storage utilization by reducing the amount of duplicate data and by sharing free and temporary space across servers
- Improve application availability
- Simplify and lower the cost of data backups through application server free backup and built in file-based FlashCopy® images
- Allow data sharing and collaboration across servers with high performance and full locking support
- Eliminate data migration during application server consolidation

- Provide a scalable and secure infrastructure for storage and data on demand IBM

TotalStorage SAN File System solution includes a Common Information Model (CIM) Agent, supporting storage management by products based on open standards for units that comply with the open standards of the Storage Network Industry Association (SNIA) Common Information Model.

Main Features

Customize Idea: Add statements or reword one or more of the following features with specific examples of your customer's SFS applications.

Direct data access by exploitation of SAN technology

SAN File System uses a data access model that allows client systems to access data directly from storage systems using a high-bandwidth SAN, without interposing servers. Direct data access helps eliminate server bottlenecks and provides the performance necessary for data-intensive applications.

Global namespace

SAN File System presents a single global namespace view of all files in the system to all of the clients, without manual, client-by-client configuration by the administrator. A file can be identified using the same path and file name, regardless of the system from which it is being accessed. The single global namespace shared directly by clients also reduces the requirement of data replication. As a result, the productivity of the administrator as well as the users accessing the data is improved.

It is possible to restrict access to the global namespace by using a non-uniform SAN File System configuration. In this way, only certain SAN File System volumes and therefore filesets will be available to each client.

File sharing

SAN File System is specifically designed to be easy to implement in virtually any operating system environment. All systems running this file system, regardless of operating system or hardware platform, potentially have uniform access to the data stored (under the global namespace) in the system. File metadata, such as last modification time, are presented to users and applications in a form that is compatible with the native file system interface of the platform.

SAN File System is also designed to allow heterogeneous file sharing among the UNIX and Windows client platforms with full locking and security capabilities. By enabling this capability, heterogeneous file sharing with SAN File System helps increase performance and flexibility.

Locking

The SAN File System is designed to provide file level locking to the SAN File System clients running on your application servers. This function provides data sharing with integrity across heterogeneous application servers. The file level locking uses the

same locking permissions (read, write, execute, and so on) that are provided by the local operating systems for locking within one server. In addition to locking the entire file, the SAN File System also supports the ability to do a lock on a byte range within a file.

Policy-based automatic placement

SAN File System is aimed at simplifying the storage resource management and reducing the total cost of ownership by the policy based automatic placement of files on appropriate storage devices. The storage administrator can define storage pools depending on specific application requirements and quality of services, and define rules based on data attributes to store the files at the appropriate storage devices automatically.

Policy-based file space provisioning

Storage administrators will be able to create rules for the SAN File System to determine what storage pool is used to allocate space for a file when the file is created. The storage administrator can use any of the file attributes (such as file name, file type, date created, user ID or group ID, and so on) to create these rules. Using the policy-based rules, the SAN File System automates the task of allocating space on the desired storage volumes.

Policy-based movement of files between storage pools and policy-based file deletion

The SAN File System provides information lifecycle management capabilities allowing administrators to specify how files should be automatically moved among storage pools during their lifetime, and optionally, specify when files should be deleted. This feature can potentially lower the overall costs of storage and improve storage space utilization enabling a balanced use of premium and inexpensive storage matching the objectives of the enterprise. This feature also has the potential to reduce manual storage administration for managing space utilization and can reduce the cost of storage management.

Lifecycle management

SAN File System provides the administrator with policy-based data management that automates the management of data stored on storage resources. Through the policy-based movement of files between storage pools and the policy-based deletion of files, there is less effort needed to update the location of files or sets of files. Free space within storage pools will be more available as potentially older files are removed. The overall cost of storage can be reduced by using this tool to manage data between high/low performing storage based on importance of the data.

Storage pools

The SAN File System provides the capability to assign storage volumes to storage pools where the storage pools can represent different classes of service for storage. Volumes can be added to a storage pool without impacting application processing. Data can be moved between volumes without affecting availability or performance.

Additional features of the IBM TotalStorage SAN File System are identified below.

Storage quota management

The SAN File System provides storage capacity to applications and servers creating new files or extending existing files. The SAN File System allows the storage administrator to define the amount of storage to be used by a specific server or by a specific application (defined as a fileset) and can help the storage administrator monitor the storage utilization by alerting them when the storage reaches a pre-defined capacity. The use of storage quotas can prevent an application from using more storage than permitted. The SAN File System allows storage utilization quotas to be set higher than the physical storage available to the system or applications. This allows for simplified management of the system by allowing administrators to dynamically add storage without having to reset application level policies. Both hard and soft quotas are supported along with threshold alerts for integration with systems management tools.

Heterogeneous file sharing between Windows and UNIX servers

The heterogeneous file sharing feature enables flexible, security-rich, sharing of files between UNIX and Windows servers. The SAN File System provides cross-environment access checking such that files created on Windows can be accessed by authorized users on UNIX, and vice versa. To control cross-environment authorization, the administrator manages a set of user map entries using the administrative CLI or GUI. Each user map entry specifies a UNIX domain-qualified user and a Windows domain-qualified user that are to be treated as equivalent for the purpose of checking file access permission in the cross-platform case. The SAN File System metadata server cluster accesses the customer's UNIX and Windows directory services, as needed, to obtain user ID and group membership information.

Caching

The SAN File System clients cache the metadata that they get from the SAN File System metadata servers. This allows the applications running on a server to continue to access the file data without additional requests to the metadata servers and allows other applications on the same server to share files without additional requests to the metadata servers. The caching of the metadata provides for high performance data access over the SAN without requests flowing to the metadata server. It can also provide continuous access to open files even if a temporary condition prohibits communication with the metadata servers.

Volume drain

The SAN File System provides a volume drain function that allows you to remove or upgrade storage volumes without disrupting application processing. The volume drain function moves the data off the specified volume onto other volumes in the same storage pool. The SAN File System manages the access to the data that is being moved so there is no need to stop application processing during the volume drain. Data can be moved between volumes without affecting availability or performance.

Flashcopy

The SAN File System provides a file-based point-in-time copy function called FlashCopy. The FlashCopy image is a logical point-in-time image of a Fileset, a set of files. A copy of the data block is made upon the next update after the FlashCopy image. The FlashCopy image only requires storage for changed blocks, therefore, reducing the amount of storage needed. The SAN File System supports 32 FlashCopy images for each Fileset. The files can be restored to any one of the FlashCopy images. File-based backup and restore products can back up files using the FlashCopy function while the applications are online and updating a later version of the files. The file-based FlashCopy function can help significantly reduce maintenance windows for data backups and can help reduce the amount of storage space needed for on-line data backups.

File copy migration tool

The SAN File System provides a file copy migration tool that you can use to copy files into the SAN File System. The file copy migration tool runs on your application servers that are the SAN File System clients. This tool includes a planning option that creates an estimate of the time needed to copy the files into the SAN File System.

iSCSI attached clients and storage

The SAN File System supports host application servers and storage devices attached through iSCSI. The iSCSI attached storage devices must conform to the SCSI driver interface standards, including the unique device identification calls.

SAN File System architecture

Customize Idea: replace Figure 1 with an appropriately labeled chart using your customer's applicable operating systems. Edit the text with your customer's application and database server names.

SAN File System architecture and components are illustrated in Figure 1 below. Computers that want to share data and have their storage centrally managed are all connected to the SAN. In SAN File System terms, these are known as clients, since they access SAN File System services, although in the enterprise context, they would most likely be, for example, database servers, application servers, or file servers.

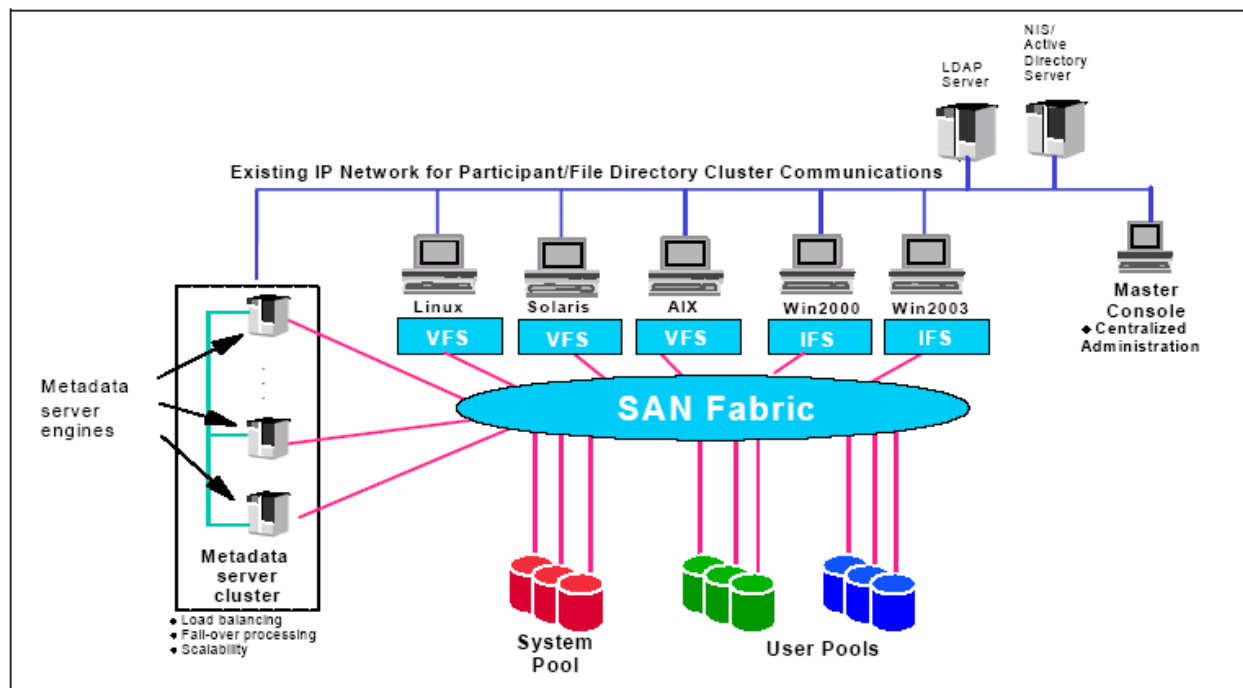


Figure 1 - SAN File System architecture

In Figure 1, we show five such clients, each running a SAN File System currently supported client operating system. The SAN File System client software enables them to access the global namespace through a virtual file system (VFS) on UNIX/Linux® systems and an installable file system (IFS) on Windows systems. This layer (VFS/IFS) is built by the OS vendors for use specifically for special-purpose or newer file systems.

There are also special computers called Metadata server (MDS) engines that run the Metadata server software, as shown on the left side of the figure. The Metadata servers manage file system metadata (including file creation time, file security information, file location information, and so on), but the user data accessed over the SAN by the clients does not pass through the Metadata servers.

This virtually eliminates the performance bottleneck from which many existing shared file system approaches suffer, giving near-local file system performance. MDSs are clustered for scalability and availability of metadata operations and are often referred to as the Metadata server cluster. In a SAN File System server cluster, there is one master Metadata server and one or more subordinate Metadata servers. Each MDS runs on a separate physical engine in the cluster. Additional Metadata servers can be added as required if the workload grows, providing solution scalability.

Storage volumes that store the SAN File System clients' user data (User Pools) are separated from storage volumes that store metadata (System Pool), as shown in the figure.

The Administrative server allows SAN File System to be remotely monitored and controlled through a Web-based user interface, called the SAN File System console. The Administrative server also processes requests issued from an administrative

command line interface (CLI), which can also be accessed remotely. This means the SAN File System can be administered from almost any system with suitable TCP/IP connectivity. The Administrative server uses an LDAP server (customer supplied) to look up authentication and authorization information about the administrative users. The primary Administrative server runs on the same engine as the master Metadata server. It receives all requests issued by administrators and communicates with Administrative servers that run on each additional server in the cluster to perform routine requests.

Example OF IBM SAN File System Exploitation - Data Backup and Recovery

Because SAN File System is a global namespace (the files are visible to all clients), the files can be backed up from any SAN File System client. Therefore, you can back up those files, either directly from the filesets or from a SAN File FlashCopy image on a completely separate SAN File System client from the client that normally runs any applications on these files, thus giving application server-free backup. This eliminates the application servers themselves from the data path of the backup and frees them from expending any CPU cycles on the backup process. If you back up the files from a FlashCopy image, this effectively almost eliminates the “backup window,” that is, a period of outage of the application to clients, since you create an online consistent copy of the data that is then backed up. The application then proceeds uninterrupted while the backup is executed against the FlashCopy image.

This principle is shown in Figure 2.




Figure 2 - Exploitation of SAN File System with ITSM

IBM SAN File System Components

IBM TotalStorage SAN File System Metadata Server Hardware and Software Requirements

<http://www-1.ibm.com/support/docview.wss?rs=575&uid=ssg1S1002518>

IBM TotalStorage SAN File System Supported Host Systems

<http://www-1.ibm.com/support/docview.wss?rs=575&uid=ssg1S1002521>

IBM TotalStorage SAN File System Supported Storage Devices

<http://www-1.ibm.com/support/docview.wss?rs=575&uid=ssg1S1002520>

IBM SAN File System Administrative Interface

There are two methods for managing the SAN File System, a command-line interface (CLI) or a graphical user interface (GUI), which is called the SAN File System console. The CLI is accessed either by logging in directly to the engine itself (at the KVM) or by using a Secure Shell (SSH) client to remotely connect to the engine. The console is accessed using a Web browser. SAN File System provides you with different levels of user access to perform administrative tasks. The user roles are defined on your LDAP server. Therefore, a user cannot be granted access to the SAN File System without using a valid user that is defined on the LDAP server.

IBM SAN File System Security

A Lightweight Directory Access Protocol (LDAP) server is used for authentication with the SAN File System administration interface. This LDAP server can be any compliant implementation, running on any supported operating system.

The SAN File System administrative infrastructure is designed to use the Secure Sockets Layer (SSL) as part of its web server software-platform base. Operators log in to the SAN File System console, which is the web-based user interface, through an SSL-secured login, and the operator ID is authenticated using the LDAP server. To use the administrative command-line interface, operators must first log in to an engine and be authorized, using the standard UNIX security mechanisms. Security can be set so that SSH-based logins to the cluster are required before the user can use the administrative command-line interface. Then, while using the administrative command-line interface, all access is authorized using the LDAP server. After authenticating the user ID, the administrative server interacts with the metadata server to process the request.

SAN File System provides different levels of user access that are assigned to specific administrative tasks in your environment. These access levels, or *user roles*, are one way to provide security. The following table describes the SAN File System user roles.

Role	Level	Description
Monitor	Basic level of access	Allows you to obtain basic status information about the cluster, display the message logs, display the rules in a policy, and list information regarding SAN File System elements such as storage pools, volumes, and filesets.
Backup	Monitor + backup access	Allows you to perform backup and recovery tasks in addition to all operations available to the Monitor role.
Operator	Backup + additional access	Allows you perform day-to-day operations and tasks requiring frequent modifications, in addition to all operations available to the Backup and Monitor roles.
Administrator	Full access	Provides you with full, unrestricted access to all administrative operations.

Table 1. SAN File System user roles

Services for IBM TotalStorage SAN File System

The following are examples of services available from IBM which can help speed the implementation of your SAN File System.

STG Implementation Services for SAN File System

- Requirements gathering & planning
- Knowledge transfer for up to two customer technical personnel
- Install and configure up to two customer supplied metadata servers including SuSE Linux®
- Install and configure IBM® TotalStorage® SAN File System (SFS) software
- Install SFS client software and assist with connection of up to five client systems
- Assist with the configuration of up to two backend storage controllers, one for metadata and one for user data with up to three LUNs each
- Configure up to two storage pools and two filesets and perform verification testing
- Optional: Provide LDAP implementation service using OpenLDAP or IBM Directory Services
- Optional: Install, configure and enable one customer supplied windows server and Master Console Software

STG Migration Services for SAN File System (Non-disruptive data migration)

- Requirements gathering & planning
- Install and configure migration software (Manager and Server Engine) on a minimum of two SFS clients
- Perform the optional task of installing and configuring the migration software (Manager only) on a stand-alone host with TCP-connectivity to the aforementioned SFS clients
- Offer customer multiple migration scenarios to be performed, such as: (1) JFS to SFS (for AIX® SFS clients), (2) NTFS to SAN File System (for Windows SFS clients). These migration scenarios shall consist of either a file synchronization or block data synchronization
- Perform migrations of JFS (AIX) data or NTFS (Windows) data to an SFS Client for either of the following applications: (1) Microsoft SQL Server, (2) Microsoft Exchange, (3) Oracle, (4) IBM DB2®, or (5) File Server
- Perform application verification testing on the migration process and analyze the migration performance

- Knowledge transfer for up to two customer technical personnel

For more information on IBM Software Services for Tivoli including contact information, please visit:

<http://www-306.ibm.com/software/tivoli/services/contact/contacting.html>

Pricing Summary

Customize this table to summarize your pricing information.

Component	Description	List Price

Insert any terms and conditions here.



Disclosure Statement

THESE DISCLAIMERS WILL NOT BE INCLUDED IN BP QP VERSION.

The information in this proposal shall not be disclosed outside the **CUSTOMER** organization and shall not be duplicated, used or disclosed in whole or in part for any purpose other than to evaluate the proposal, provided that if a contract is awarded to IBM as a result of or in connection with the submission of this proposal, **CUSTOMER** shall have the right to duplicate, use or disclose the information to the extent provided in the contract. This restriction does not limit the right of **CUSTOMER** to use information contained in the proposal if it is obtained from another source without restriction.

This proposal is not an offer or contract. Neither IBM nor you have any obligations or liability to the other unless our authorized representatives enter into definitive written agreement. Terms included in this proposal are not binding unless they are included in such a written agreement.

This proposal is valid for 30 days unless otherwise noted. Prices and charges in this proposal are subject to change without notice and do not include applicable sales taxes. The prices quoted here reflect IBM's current prices as of the date of quote, for the applicable model/ configuration. IBM products are subject to withdrawal from marketing and or service upon notice, and changes to product configurations, or follow-on products, may result in price changes. This document and all information herein is provided AS IS, without warranty, and nothing herein, in whole or in part, shall be deemed to constitute a warranty. IBM Machines are subject to the Statements of Limited Warranty accompanying the applicable Machine. IBM Program Products are provided subject to their applicable license terms.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM is not responsible for printing errors in this proposal that result in pricing or information inaccuracies. Products, programs, services or features discussed in this proposal may be subject to change without notice.

IBM hardware products may be manufactured from new parts, or new and used parts. Regardless, IBM warranty terms apply.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country. This information may be subject to change without notice.

DELETE IBM DISCLAIMERS FOR BP VERSION ONLY TO THIS POINT

Trademarks

IBM, the IBM logo, AIX, DB2, FlashCopy, Tivoli, and TotalStorage are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both.

Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

The IBM home page can be found at: <http://www.ibm.com>