HITACHI
DATA SYSTEMS

A New Look at Remote Data Replication for the Small and Medium-Sized Enterprise

Reducing the Cost and Complexity of Remote Synchronous Data Replication Over a Wide-Area IP Network

A White Paper

By Terry Buchanan, CONPUTE;
with Thomas Neidhardt and
Deborah Howard, McDATA;
Lee Lipes, Redline;
and Christophe Bertrand,
Hitachi Data Systems

August 2006

CONPUTE

McDATA®

Redline communications

Executive Summary

Real-time replication of critical information to a remote data center offers an organization the most effective safeguard against data loss from an unforeseen outage. Unlike tape-based data protection strategies, with which recovery time is often measured in days, synchronous remote replication allows users to resume accessing business applications in minutes, with little, if any, data loss. The one caveat with wider use of synchronous data replication, however, is the high cost of telecommunications networks.

According to IDC, telecommunications expenses represent 50 percent to 70 percent of the total ongoing cost of implementing data replication. The requirement for fast network connectivity between local and remote data centers places a significant load on an organization's IT budget. A T3 connection, offering 45 megabits per second (Mbit/sec) of bandwidth, is considered the minimum configuration, and large enterprises, with significant amounts of critical operational data, can require substantially bigger pipes. With hefty one-time installation charges and ongoing fees that can easily surpass US$15,000 per month, the minimum network connectivity requirement for replication has a high entry price.

Midrange storage systems now routinely support enterprise-class data replication functionality, holding out the tantalizing promise of improved business continuity planning capabilities for small to medium-sized enterprises (SMEs). However, the cost of providing network connectivity between a local data center and a remote one represents a significant barrier to broader replication deployment.

Now, through a novel combination of existing off-the-shelf technologies, the telecommunication costs associated with remote data replication have been dramatically reduced. Instead of using costly traditional telecommunications connectivity, this new solution leverages existing products that provide protocol-converting storage routing and fixed wireless IP technology. These two components allow existing Fibre Channel storage systems to cost-effectively pass replication traffic between remote and local data centers at speeds equivalent to those of T3s, without relying on traditional telecommunications infrastructure cabling. With a one-time fixed entry price for communication between two buildings—equivalent to a few months of T3 access charges—fixed wireless IP technology removes a significant economic barrier, giving SMEs access to the same remote replication capabilities large enterprises enjoy.

Hitachi Data Systems, McDATA, and Redline Communications are partnering to bring SMEs a unique approach to remote data replication. Using the powerful remote replication capabilities of the Hitachi TrueCopy™ Synchronous software, the protocol conversion and compression features of the McDATA Eclipse 1620 storage router, and the broadband fixed wireless technology offered by Redline Communications, SMEs can now implement metro-area remote data replication for a fraction of the cost of a traditional landline-based implementation.

Designed in response to real-world customer problems, the visionary combination of Hitachi storage, McDATA storage area network (SAN) routing hardware, and Redline communications technology was developed and is supported by CONPUTE, a Gold-level partner and an Installation and Configuration Certified partner in the Hitachi TrueNorth™ Channel Partner Program. "This solution was a direct response to an assessed client need for lower-cost remote data replication," says Terry Buchanan, vice president of services at CONPUTE. "We are seeing the same issue over and over again at many of our clients. The high cost of telecommunications data lines is preventing the wider deployment of remote replication and compromising more-effective business continuity planning for SMEs." An acknowledged expert in IT disaster recovery best practices, Buchanan provides consulting to CONPUTE clients on strategic business continuity planning and risk evaluation and control.

Contents

A New Look at Remote Data Replication for the Small and Medium-Sized Enterprise

Reducing the Cost and Complexity of Remote Synchronous Data Replication Over a Wide-Area IP Network

Business Continuity Planning for Small and Medium-Sized Enterprises

Contingency planning has come under increased scrutiny from executives, investors, regulators, and customers, as awareness has grown of the catastrophic impact a system outage can have on an organization's ability to conduct business. With day-to-day operations increasingly dependent on the IT infrastructure, IT managers of even the smallest enterprise are under pressure to maintain business continuity plans that promise faster resumption of business and no loss of data following an outage.

Preparing for unforeseen application downtime is now a legitimate cost of doing business. The risk of revenue loss from an outage is high, with the potential to have an impact on corporate valuation and customer credibility. In addition, for many industries, regulatory oversight from government agencies mandates data protection requirements, with hefty fines and criminal penalties for noncompliance.

With so much riding on the availability of corporations' data assets, approaches to business continuity planning have grown in sophistication. What was once a one-size-fits-all process—with all applications receiving the same tape-based data protection treatment—has now evolved, allowing higher-value data to receive priority.

Figure 1. The Cost of Business Continuity

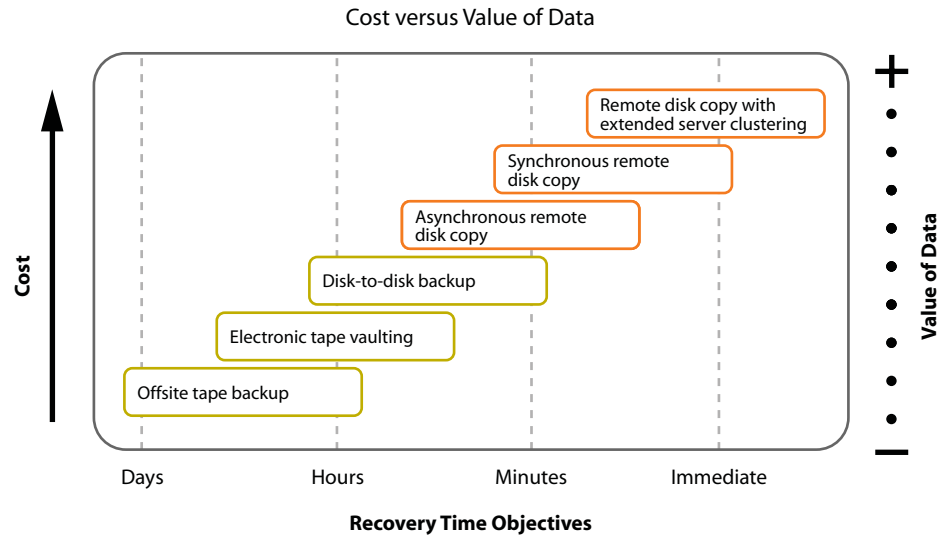

Enterprises can now build business continuity plans based on the value of their data, with cost of data protection (from offsite tape backup to remote disk copy with extended server clustering) tied directly to the value and/or the accessibility requirements of the data itself.

Remote data replication is the technology of choice for organizations with business-critical applications that cannot afford to be offline or suffer data loss. Data replication technology mirrors every application write request to a storage system hosted in a distant data center. By keeping two copies of critical data perfectly synchronized, data replication technology provides redundancy, allowing the organization to weather a system failure. If one copy of the data is lost during an outage, business applications can be quickly switched to the alternate copy without suffering loss of data or lengthy downtime.

Synchronous Data Replication

Synchronous replication involves two storage volumes hosted on a network capable of carrying data traffic—usually interconnected Fibre Channel SANs. When an application initiates a request to modify data on the primary storage volume, the replication software sends an identical I/O request to the remote storage volume. The initiating write request is not considered complete until both I/Os—local and remote—are successful. If anything happens to cause one I/O to fail, both I/Os are rolled back. This ensures that the remote storage volume is always an exact copy of the local volume.

Figure 2. Synchronous Remote Replication at the LUN Level

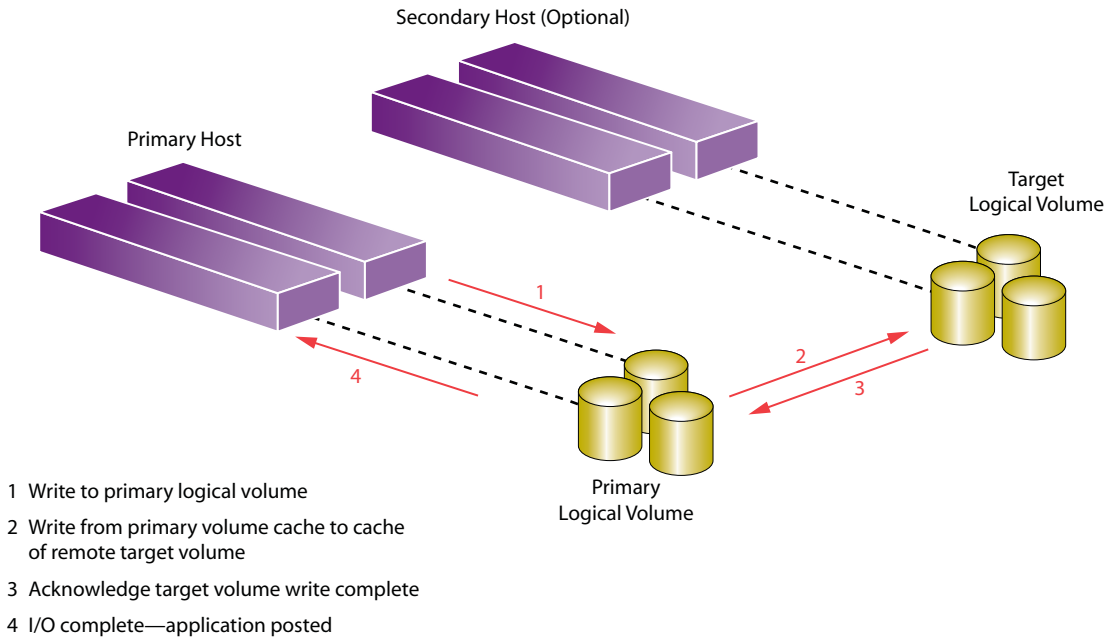

The tight synchronization provided by synchronous remote replication is ideal for applications that cannot tolerate loss of data and that require fast recovery. The one drawback of synchronous replication is the need for geographical proximity.

Distance Limitation

During synchronous remote replication, local and remote I/Os are tightly coupled. Any latency in the network telegraphs directly back to the business application as slower performance. Although Fibre Channel SANs are capable of spanning hundreds of kilometers, performance considerations set a practical limit for synchronous replication of 35 to 50 kilometers (km). At this distance, propagation delays have minimal impact on the application.

Although the distance limitation prevents synchronous replication from being used when data centers are widely dispersed geographically, it does allow for replication within metro areas and provides a significant improvement over business continuity planning with tape-based data protection. The distance limitation also opens up the options for network connectivity, with the potential to significantly reduce telecommunications costs.

The Cost of Network Connectivity

Bandwidth is a significant recurring operational cost of any long-distance data replication configuration. And although synchronous replication provides the fastest means of recovering applications after an outage, the technology is rarely deployed outside of large enterprises, due to the high cost of network connectivity. A T3 data line, providing 45 megabits per second (Mbit/sec) throughput, is considered the minimum necessary connectivity for data replication. Most large enterprises demand significantly more bandwidth, requiring OC3, OC12, OC48, or OC192 data lines, which provide connectivity of anywhere from 155Mbit/sec (OC3) to 9.6 gigabits per second (Gbit/sec) (OC192).

Although the pricing and availability of data lines vary considerably by location, a T3 line is likely to cost between US$15,000 and US$21,000 per month. The enterprise-targeted connectivity of OC3, OC12, OC48, or OC192 data lines requires a substantial up-front installation investment and recurring monthly charges, based on the distance between endpoints. Prices for these high-bandwidth connectivity options start at the high end of T3 pricing and escalate rapidly.

For SMEs operating on a tight IT budget, the ongoing monthly expense of a T3 line represents a huge economic hurdle to remote replication deployment. However, since synchronous data replication is in essence restricted to network connectivity spanning no more than 50km, another networking option is needed. This alternative technology proves to be an order of magnitude simpler and more cost-effective than traditional telecommunications landline connectivity.

Broadband Fixed Wireless Networks

Fixed wireless point-to-point microwave technology has been in use for a long time and has traditionally been the choice of telecommunications carriers for long-distance backhaul networks. Broadband fixed wireless (BFW) networks, an evolution of the traditional fixed wireless technology, operate in the license-exempt (LE) radio spectrum, allowing solutions to be deployed without registration with the U.S. Federal Communications Commission (FCC). BFW solutions offer high-speed point-to-point connectivity over distances of up to 80km, with wider coverage available with multiple hops. These networks support data traffic rates of up to 48Mbit/sec, surpassing the capacity of a land-based T3 connection.

For SMEs looking to deploy business continuity solutions based on synchronous data replication, carrier-grade BFW networks offer a persuasive alternative to the traditional telecommunications infrastructure. With no long-distance cabling requirements, a low-tech implementation that can be handled by nontechnical staff, and a one-time fixed capital cost for deployment, BFW networks dramatically extend the capabilities of midrange storage system–based data replication without significant impact on the IT budget.

Elements of a BFW Data Replication Configuration

Although BFW networks and data replication solutions are not new, it is the innovative combination of off-the-shelf products using these technologies that provides a compelling business continuity approach for SMEs. Hitachi Data Systems, McDATA, and Redline Communications have partnered to bring these technologies together for SMEs for the first time. The solution provides enterprise-class business continuity at a significantly lower cost than that of traditional implementations.

Figure 3. Elements of a BFW Data Replication Configuration

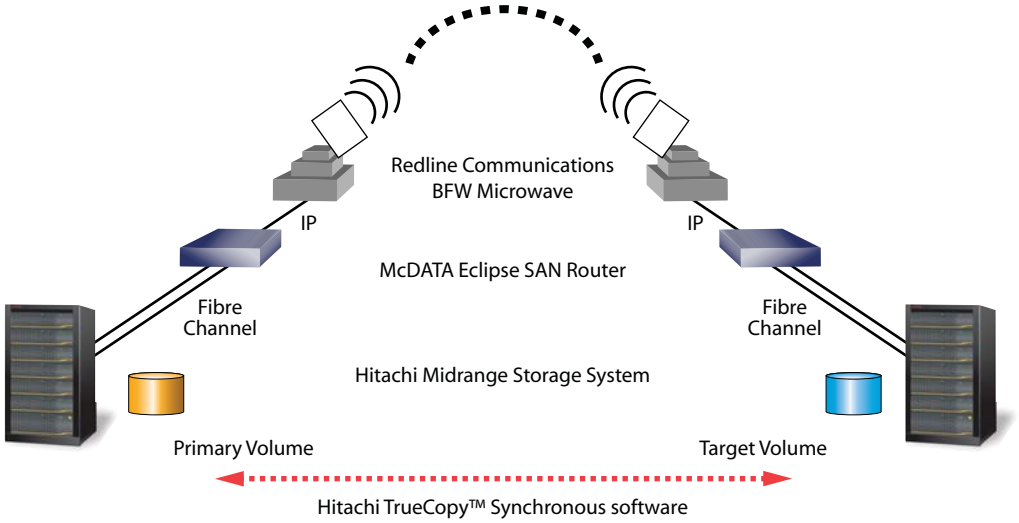

Hitachi Data Systems, McDATA, and Redline Communications have partnered to bring these technologies together for SMEs for the first time, offering enterprise-class business continuity at a significantly lower cost than that of traditional implementations.

Hitachi Midrange Storage Systems and Hitachi Remote Replication Software

The midrange Hitachi TagmaStore® Adaptable Modular Storage systems provide cost-effective, highly available, high-performing storage solutions for SMEs and large enterprises. The versatile and scalable storage systems support a wide range of advanced storage functionality, the intermix of Fibre Channel and SATA disk drives in the same system, and "single pane of glass" management.

TrueCopy Synchronous software is an innovative, storage-based replication solution that operates on Hitachi midrange storage systems. It provides synchronous data replication capabilities. The redundant copies of critical business data it creates can be used for rapid recovery of business applications after an outage.

Operating on most new and legacy Hitachi midrange storage systems, TrueCopy Synchronous software provides a flexible and cost-effective remote copy solution. Replication is performed on the storage system controller, eliminating processing overhead from the application host server. The primary and secondary replication

volumes do not have to be of the same type or RAID configuration. For example, primary volumes in the main data center can be high-speed Fibre Channel devices with a RAID configuration tuned for performance, and secondary volumes at the remote data center can be cost-effective storage with RAID optimized for efficient use of capacity.

Although the Hitachi midrange storage systems operate on a Fibre Channel host interface, the use of protocol-converting storage routers increases the connectivity options, allowing flexible IP-based BFW network connectivity without affecting performance.

McDATA SAN Routers

The McDATA Eclipse 1620 SAN Router provides a bridge between the Fibre Channel fabric, hosting the Hitachi midrange storage systems, and the extended IP-based BFW network. With support for a wide variety of standards-based communication protocols—Fibre Channel, Ethernet (Fast Ethernet and Gigabit Ethernet), iSCSI, and iFCP—Eclipse routers allow two Fibre Channel-based Hitachi storage systems, separated by an Ethernet network, to perform data replication.

Synchronous data replication is sensitive to the speed of the network, and any latency is immediately reflected back to the business application as poor performance. The Eclipse 1620 is a true hardware-based router, offering low latency and high throughput to maximize I/O per second (IOPS). A simple, intuitive management interface, accessed through an IP connection, allows the router to be monitored and managed from a central location. And various dynamically tunable configuration options help minimize latency and maintain wire-speed communication at all times.

Compression

Compression can significantly improve use of the available network bandwidth. The SAN router employs configurable intelligent compression routines that dynamically optimize the traffic on the Ethernet port to ensure the fastest communication possible. The compression routines look for repetitive patterns in the data stream as it passes across the intelligent ports of the router. When set to automatic, the router dynamically adjusts the amount of data compression applied to allow the data rate at the Ethernet exit port to be as close to optimal as possible. For example, if uncompressed data is arriving at 65Mbit/sec and the port speed of the Ethernet port is set for T3 (approximately 45Mbit/sec), the data stream will be compressed just enough to ensure that T3 speeds are maintained. An alternative compress-always mode, meant for use with a shared network connection, ensures that the minimum-possible bandwidth is used.

The compression method—LZO, Fast LZO with History, LZO with History, or Deflate options—is configured at the router. A new high-speed, low-latency hardware mode is also available, providing full-wire-speed, full-duplex compression on networks supporting gigabit connectivity. The various user-configurable modes allow router compression to be optimized to achieve a balance between speed and network capacity. Reports generated by the router's Element Manager software show the success of a chosen compression method.

Jumbo Frames

The McDATA SAN router optimizes the conversion of Fibre Channel data packets for transmission over an Ethernet network. Under normal circumstances, two Ethernet frames are required to carry the payload from a single Fibre Channel frame. The SAN router uses Jumbo Frames to extend the Ethernet payload capacity.

Jumbo Frames processing treats the Fibre Channel frames from a session as a single data stream. Before mapping the stream to Ethernet packets, it applies data compression and factors in the maximum transmission unit (MTU)—the Ethernet frame size. This process ensures that a single Fibre Channel frame is not split into two Ethernet packets. It also optimizes communications, reducing network traffic and latency.

Fast Write

The McDATA SAN router also employs the patent-pending Fast Write capability, which allows wire-speed gigabit performance to be maintained over any distance. The SCSI protocol at the heart of all Fibre Channel network communications was initially designed with the assumption that the host and storage would be in close proximity. When performing write operations, the SCSI host and storage send multiple handshakes and acknowledgments back and forth to ensure the integrity of the data transfer. However, with today's storage networks, capable of spanning hundreds—if not thousands—of miles, the complex series of SCSI handshakes and acknowledgments introduces unnecessary network traffic and unwanted network latency.

Figure 4. McDATA Fast Write Optimized SCSI Write

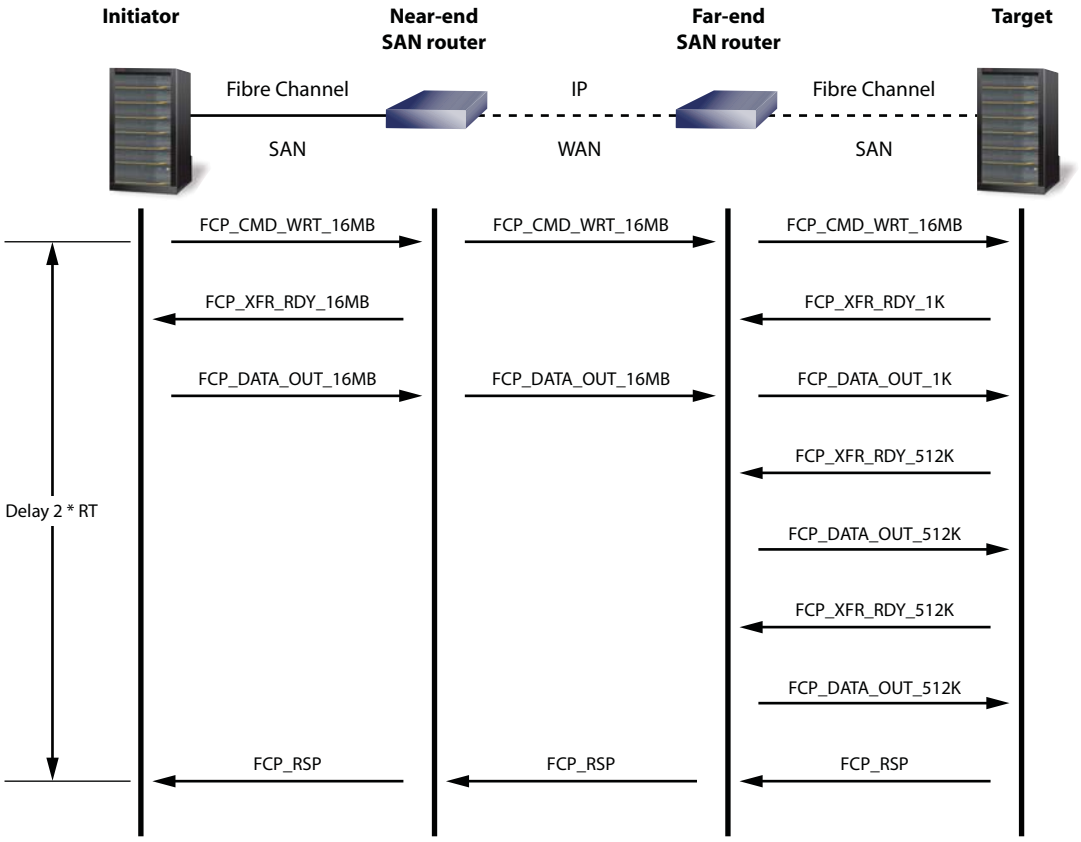

Fast Write allows wire-speed gigabit performance to be maintained over any distance, eliminating much of the time-consuming back-and-forth communication in a traditional SCSI write operation.

Fast Write eliminates much of the time-consuming back-and-forth communication in a traditional SCSI write operation. Implemented in the hardware of the router's intelligent ports, Fast Write reduces unnecessary round-trips, speeding network communications. When a SCSI write operation is detected, the two SAN routers, acting as bridges to the extended IP network, mimic the recipient and the sender of the SCSI commands. The bandwidth-consuming handshakes and acknowledgements, which would otherwise be passed across the wide-area network, are performed locally, with only the data and final-command-completion message passing over the wider network infrastructure. Fast Write speeds up SCSI processing, eliminating network traffic and latency, even when only a short distance separates the two routers.

Port Priority

The McDATA SAN router increases the flexibility of a replication configuration, with support for port-based priority. This feature allows replication traffic to be classified according to importance, with data from critical business applications receiving a higher priority on the network.

Redline Communications Broadband Fixed Wireless

Redline Communications is a provider of BFW systems that operate in the LE radio spectrum. LE-band solutions communicate at 5.725GHz to 5.825GHz and can be implemented without regulation or oversight from the FCC. This allows BFW networks using the LE band to be deployed quickly and at relatively low cost.

Low Cost

It is estimated that between 50 percent and 70 percent of the total cost of a data replication implementation derives from network connectivity. In a conventional replication configuration, two or more data centers are connected by a high-speed land-based telecommunications network—T3, OC3, or beyond. For organizations in a location already served by a telecommunications carrier, a land-based data line represents a significant ongoing monthly investment—a T3 connection averages US$15,000 per month, and an OC3 implementation carries installation costs, a minimum monthly charge, and variable costs based on the distance between the two locations being connected. For businesses in a location with no high-speed data telecommunications carrier, the cost of laying new cable to connect to an existing infrastructure can represent an insurmountable hurdle.

In contrast, a BFW network, based on the Redline AN-50, provides T3 communication speeds across metro-area distances, for a fixed one-time cost. The difference between the cost of traditional telecommunications carrier bandwidth and that of a Redline Communications BFW network is enormous.

No Long-Distance Cabling

The Redline AN-50 requires no special cabling to deploy a T3-comparable network. Because communications are broadcast in the LE radio frequency, the solution is ideal for organizations in regions not served by high-speed data networks.

The solution consists of a small dish antenna and a transceiver containing an Ethernet port. The microwave antennae are mounted on the outside of each building or on a mast and pointed in each other's general direction. An audible signal indicates when the antennae are optimally aligned, and once positioned, the AN-50 immediately forms an Ethernet network.

No Interference

Historically, one of the major impediments to microwave networks has been interference from other users sharing the same frequency—in-band—and from transmissions in adjacent licensed frequencies—out-of-band. With users free to implement microwave solutions in the LE spectrum, without permission or consideration for others, the possibility of overlapping signals interfering with transmissions is real. The Redline BFW technology is designed with inherent tolerance for interference and uses specific patented approaches to guarantee the integrity of the microwave signal.

To combat in-band interference, the Redline AN-50 uses a broad range of techniques, including orthogonal frequency-division multiplexing (OFDM) with adaptive coding, multiple channel and time division duplex (TDD), narrow-beam antennae with high side and back lobe rejection, cross-polarization, automatic repeat request (ARQ) correction, and adaptive modulation.

Advanced OFDM with adaptive coding is the physical layer foundation of the Redline solution. This technology provides redundancy by spreading the same signal content over time and across multiple subcarriers. If interference arises and affects some of the subcarriers, others will contain sufficient content to allow the signal to be recovered. The technique is also effective against multipath fading, another common problem with microwave communication. Quick bursts of interference can also be addressed by use of ARQ error correction. Operating in the physical layer, the ARQ error correction requests the retransmission of signal packets that have been compromised. This mechanism is similar to TCP error correction, but it has much lower overhead because it operates in the physical layer. The AN-50 also uses adaptive modulation and operates on multiple channels to further mitigate the effects of interference and provide maximum flexibility for planning an installation.

The antennae used with the Redline AN-50 are designed to eliminate interference from systems operating in the same region. The combination of cross-polarization, narrow-beam antennae, and multiple channels provides an effective strategy for combating interference from co-located systems. Redline antennae are frequently installed on poles and masts that host antennae for other microwave solutions, both LE and licensed, without any problems from interference.

CONPUTE Solution Services

CONPUTE, a Gold-level partner in the Hitachi TrueNorth Channel Partner Program, pioneered the BFW networking approach to data replication, using TrueCopy software from Hitachi Data Systems and protocol-converting SAN routers from McDATA. During discussions with its clients, CONPUTE confirmed that the main obstacle to wider deployment of remote data replication was the high cost of telecommunications connectivity. With more than 30 years' experience in designing IT solutions for SME clients, CONPUTE immediately recognized the need for an alternative to traditional leased-data-line connectivity and set out to find a solution.

With an in-depth understanding of SME IT infrastructures and specialized knowledge in outsourcing, security, storage, and wireless broadband, CONPUTE devised a remote replication solution to satisfy the specific needs of its clients. Using off-the-shelf technologies, CONPUTE combined the world-class remote replication capabilities of TrueCopy software, the high-speed protocol-converting functionality of the McDATA SAN router, and the cost-effective connectivity of the Redline Communications BFW solution. The result was a simple-to-maintain, low-cost remote replication alternative for SMEs.

BFW-based Data Replication in the Real World

A CONPUTE client maintained a primary data center at its head office and a second, backup data center within the same metropolitan area. For business continuity purposes, tape backups were transported regularly between the primary and backup data centers. However, disaster-preparedness tests had produced disappointing results: application recovery was slow and always involved some degree of data loss.

Data replication was evaluated as an alternative business continuity solution, but as the cost of connectivity between the two data centers was significant, this approach was ruled out. After discussing the problem with CONPUTE, the client decided to implement the following solution.

Figure 5. BFW-Based Replication for the SME

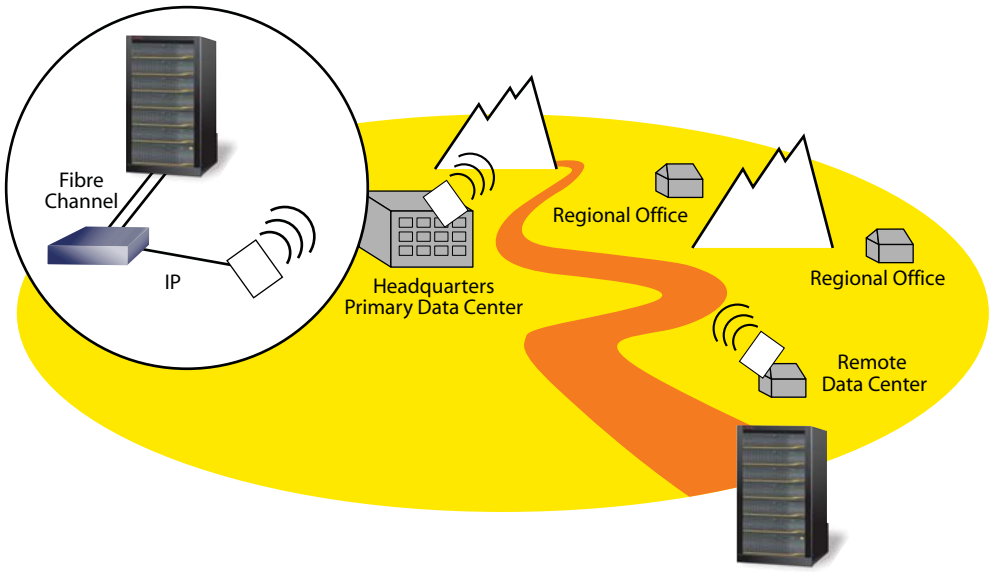

A BFW-based replication solution allowed the SME to implement network connectivity between the primary and remote data centers for a one-time capital expense equivalent to a few months of traditional telecommunications carrier T3 leasing costs.

Hitachi modular storage systems*, equipped with TrueCopy Synchronous software, were deployed at the primary and backup data centers. Each data center also received a Redline Communications AN-50 BFW microwave network antenna and transceiver and a McDATA Eclipse 1620 SAN Router. This configuration enabled the primary data center to replicate synchronously to the remote backup data center, creating a real-time redundant copy of application information for use in the event of a disaster.

In testing, CONPUTE demonstrated that 150 IOPS of changes at the primary data center generated 9.16 megabytes per second (MB/sec) of raw data. This information was then compressed to 3MB/sec by the McDATA Eclipse router and sent across the Redline AN-50 BFW microwave network to the backup data center. The 3MB/sec transmission corresponded to 24Mbit/sec network bandwidth, equivalent to 16 T1 lines, or partial use of a T3.

The BFW-based replication solution allowed CONPUTE's client to implement network connectivity between the primary and remote data centers for a one-time capital expense equivalent to a few months of traditional telecommunications carrier T3 leasing costs. This significantly reduced the cost of remote replication and dramatically improved business continuity capabilities.

Without the high-speed data network provided by the Redline Communication solution and the McDATA SAN router, this replication configuration could not function. The use of a BFW network proved to be a huge savings over conventional telecommunications carrier landlines. Not only is the BFW network more cost-effective, but the system is also simple and fast to deploy, allowing IT staff members at the primary data center to manage all network communications and storage equipment from their central location.

* Please check with Hitachi Data Systems for models that support TrueCopy Synchronous software.

Summary

With BFW microwave technology, users of midrange storage solutions can now incorporate synchronous data replication into their ongoing business continuity planning process. BFW networks provide a cost-effective alternative to conventional telecommunications carrier-based solutions. Hitachi Data Systems, McDATA, Redline Communications, and CONPUTE have successfully implemented remote synchronous replication based on microwave network connectivity. This combination of existing off-the-shelf technologies promises to revolutionize recovery-time and recovery-point objectives for the business applications of small and medium-sized enterprises. The solution delivers the same level of data protection and disaster preparedness enjoyed by large enterprise data centers, at a fraction of the cost.

Hitachi Data Systems Corporation

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com

WHP-208-02 August 2006