



SAN for Business Continuity

How Cisco IT Achieves
Subminute Recovery Point
Objective



A Cisco on Cisco Case Study: Inside Cisco IT

Overview

- Challenge

 - Improve recovery point objective (RPO) and recovery time objective (RTO) for ERP database disaster recovery

- Solution

 - Use fibre channel over IP (FCIP) with EMC SRDF/A

- Results

 - Sub-minute RPO over a distance of 3000 miles and improved RTO

- Next Steps

 - Use FCIP to unify all Cisco IT and Engineering SANs

Background: Replication Terms

- Recovery point objective (RPO): How up-to-date the database is when restored to service
- Recovery time objective (RTO): How long it takes to bring the disaster recovery (DR) database online with the desired RPO
- Perfect RPO (zero seconds) requires synchronous replication between sites
- Because Cisco's sites are 3000 miles apart, Cisco strives for $RPO < 60$ seconds

Challenge: Improve RPO and RTO for ERP Database Disaster Recovery

- Cisco uses EMC SRDF for local database replication
- What to use for replication between production and DR databases?

Production database are in San Jose, California

DR is 3000 miles away, in Research Triangle Park (RTP), North Carolina

Challenge: Improve RPO and RTO for ERP Database Disaster Recovery (Contd.)

- EMC SRDF/S in synchronous mode creates too much latency

Over fibrechannel, every 100 miles adds 1ms latency due to speed of light

Round-trip time (RTT) latency for transferring a single write I/O = 61 milliseconds

A transaction that usually takes ten seconds might take two minutes with SRDF/S: unacceptable

Challenge: Electronic Journaling

- What it is:

 - Application-based asynchronous process to support DR

 - New transactions are first stored in redo log before being committed to database volumes

 - When redo log fills up, log is copied to create an archive log

 - Archive log is copied over the WAN to RTP for application to DR database

 - Sequential archive logs contain a map for all changes occurring in production databases between two points in time

Challenge: Electronic Journaling

- Drawbacks of previous electronic journaling DR process

High management overhead required for host-based archive-replication scripts—one for each database on each host

Need to use host CPU resources for compression, performing a network push to remote DR hosts, and applying logs

Shared host resources for DR and development: the need to change partition sizes during DR took time and increased RTO

Log management issues

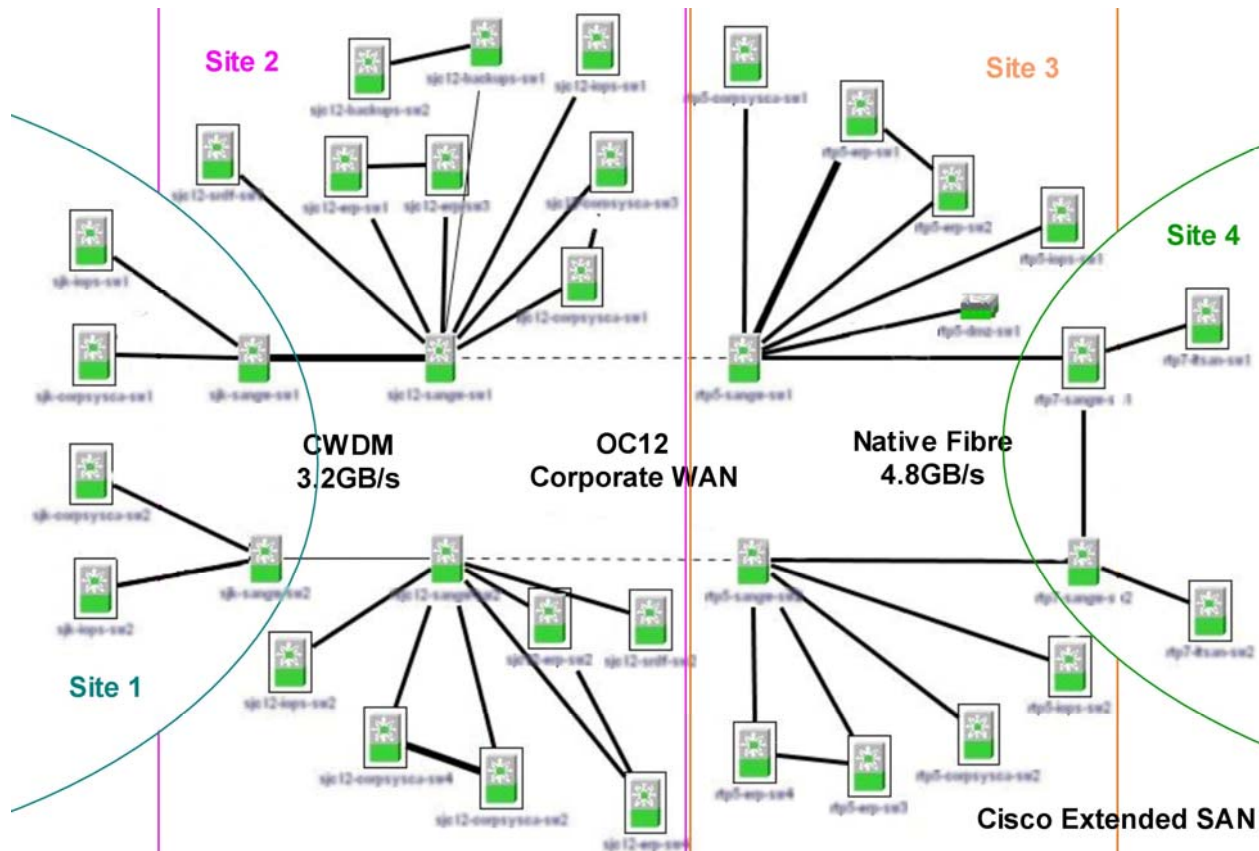
Potential for getting too far behind on applying logs to catch up

Outage required during failback

Potential for database corruption, a risk of real-time or near-time replication

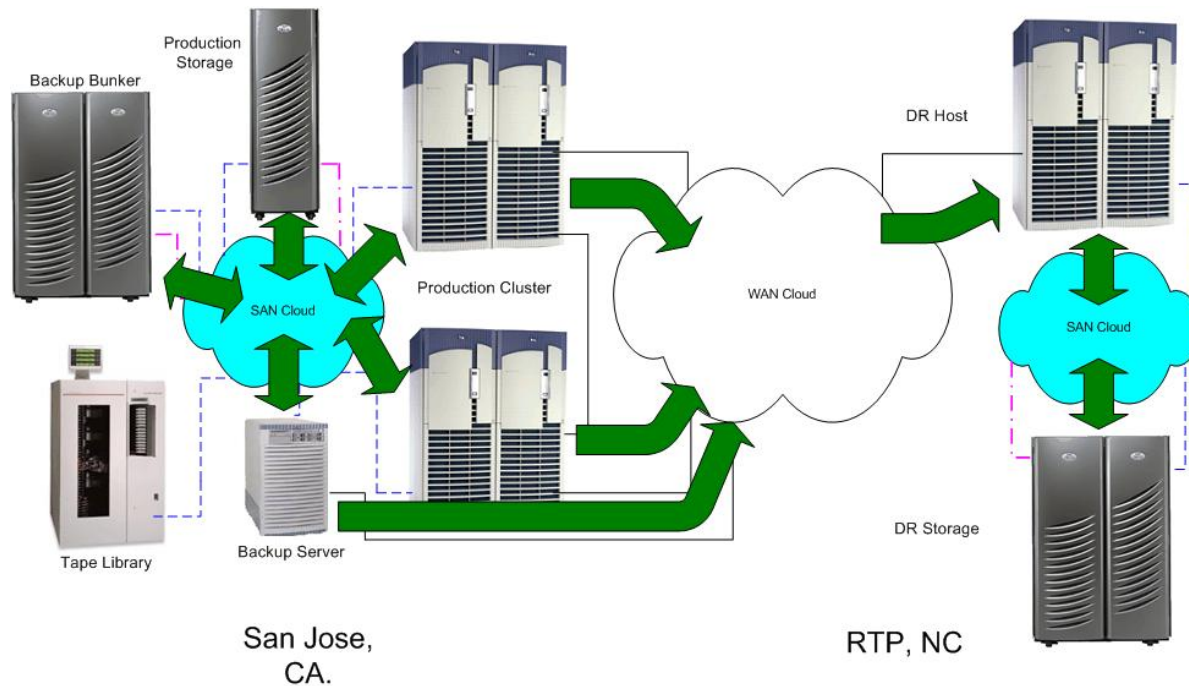
Solution: 1. Unify Cisco SANs

- Cisco connected its two Cisco MDS 9000 Family Switches, one for each fabric, with FCIP



Solution: 1. Unify Cisco SANs

BEFORE FCIP: Host-Based Replication

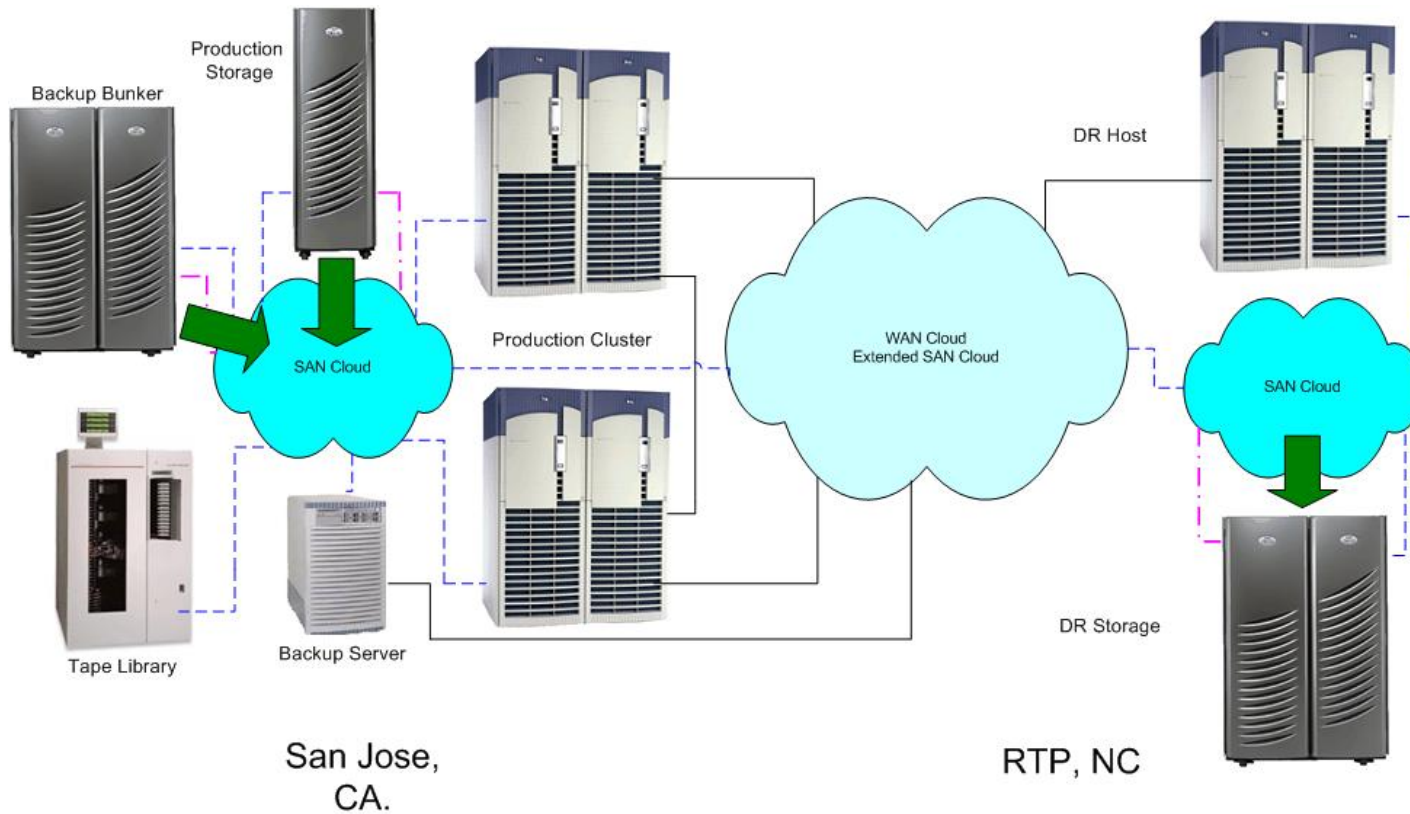


Host Based Replication Before FCIP

Green arrows show data flow

Solution: 1. Unify Cisco SANs

AFTER FCIP: Array-Based Replication



Array Based Replication With FCIP

Green arrows show data flow

Solution: 2. Achieve Incremental Relationship

- Goal: Maintain an incremental relationship between production and DR databases

If DR database falls too far behind in application of logs, refresh database from production disk backup

- Solution: Use EMC SRDF/S to copy data directly from production backup frame to DR frame over FCIP, every six to 24 hours

Only transmit changes since last backup

- Limitation: Interruption to replication process means that DR database will be inconsistent with production database, and therefore not restartable

Solution: 3. Achieve Asynchronously Updated, Consistent Relationship Without Latency

- Cisco IT conducted a pilot using EMC SRDF/A (asynchronous), designed to overcome challenges of long-distance replication
- If SRDF/A image is interrupted, no more than two cycles—60 seconds for Cisco—are lost, and remote site remains consistent and restartable
- Limitation: Each change is transferred three times, in archive log, redo log, and database volume

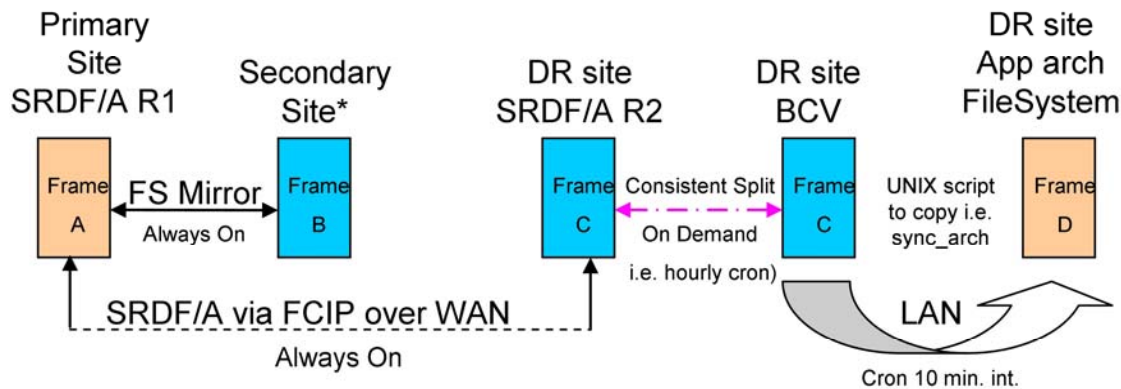
After FCIP compression, ~1.75 times the bandwidth required for write changes traveled over the WAN

Therefore, solution could not scale across multiple DR environments unless Cisco upgraded WAN to OC-48

- But, this gave Cisco IT an idea...

Solution: 4. Transfer Only the Archive and Redo Logs

By transferring only the archive and redo logs, Cisco improved RPO compared to using internally-developed scripts for archive log transfer



Archive Log Transfer Flow

Only 1-2 logs are in transit (unavailable) at any given time

Results: Greatly Improved Replication

- Sub-minute RPO over a distance of 3000 miles
- RTO measured in minutes instead of hours
- Frame-to-frame transfer rates of 50MB/second
- Incremental relationship between production and DR, enabling failback without downtime
- Ability to refresh a 4TB database over the WAN in two hours, compared to 24 hours using a host-to-host transfer

Results: Greatly Improved Replication (Contd.)

- Single SAN view of four data centers
- No need to use host resources for replication, or to assign a DR host to apply logs in standby mode
- No production impact

Next Steps: Continue to Improve Replication Strategy

- Minimize amount of data transferred across Cisco WAN
 - Establish concurrent SRDF/A relationships: one between production databases and remote DR, and another between production databases and a backup bunker
- Begin replicating Amsterdam databases in RTP, for a single view of entire IT SAN
- Use FCIP to tie Cisco IT SANs to Cisco Engineering SANs

Next Steps: Continue to Improve Replication Strategy (Contd.)

- Use FCIP as a fail-over path for campus inter-building SAN extensions, such as CWDM and dark fibre
- Implement FCIP and SRDF/A for other mission-critical ERP databases as Cisco increases WAN capacity between San Jose and RTP

To read the entire case study, or for additional Cisco IT case studies on a variety of business solutions, visit Cisco on Cisco: Inside Cisco IT

www.cisco.com/go/ciscoit



CISCO



Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883


Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 ©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0704R)