



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	PATENT NUMBER	GROUP ART UNIT	FILE WRAPPER LOCATION
10/456,053	7152182	2113	9200



Correspondence Address/Fee Address Change

The following fields have been set to Customer Number 56436 on 03/18/2016

- Correspondence Address
- Maintenance Fee Address

The address of record for Customer Number 56436 is:

56436
Hewlett Packard Enterprise
3404 E. Harmony Road
Mail Stop 79
Fort Collins, CO 80528



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	ISSUE DATE	PATENT NO.	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	12/19/2006	7152182	200309044-1	5067

22879 7590 11/29/2006
HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

ISSUE NOTIFICATION

The projected patent number and issue date are specified above.

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b) (application filed on or after May 29, 2000)

The Patent Term Adjustment is 581 day(s). Any patent to issue from the above-identified application will include an indication of the adjustment on the front page.

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) WEB site (<http://pair.uspto.gov>).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at (571)-272-4200.

APPLICANT(s) (Please see PAIR WEB site <http://pair.uspto.gov> for additional applicants):

Minwen Ji, Sunnyvale, CA;
Alistair Veitch, Mountain View, CA;
John Wilkes, Palo Alto, CA;

NOV. 2. 2006 8:28AM HP LEGAL

NO. 2899 P. 1

TO: PTO ISSUE FEES

ART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: Mail

Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

22879 7590 10/24/2006

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

Certificate of Mailing or Transmission
I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

ALBA P ESCOBAR	(Depositor's name)
<i>[Signature]</i>	(Signature)
11/2/06	(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	06/06/2003	Minwen Ji	200309044-1	5067

TITLE OF INVENTION: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

APPL. TYPE	SMALL ENTITY	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	NO	\$1400	\$300	\$0	\$1700	01/24/2007

EXAMINER	ART UNIT	CLASS-SUBCLASS
MCCARTHY, CHRISTOPHER S	2113	714-00600

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

☐ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.
☐ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. Use of a Customer Number is required.

2. For printing on the patent front page: 11/2/2006

(1) the names of up to 3 registered patent attorneys or agents OR, alternatively, 01 FC:1501
(2) the name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.111. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE

(B) RESIDENCE: (CITY and STATE OR COUNTRY)

Hewlett-Packard Development Company, LP.

Houston, Texas

Please check the appropriate assignee category or categories (will not be printed on the patent): ☐ Individual ☒ Corporation or other private group entity ☐ Government

4a. The following fee(s) are submitted:

☒ Issue Fee
☐ Publication Fee (No small entity discount permitted)
☐ Advance Order - # of Copies

4b. Payment of Fee(s): (Please first reapply any previously paid issue fee shown above)

☐ A check is enclosed.
☐ Payment by credit card. Form PTO-2038 is attached.
☒ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment, to Deposit Account Number 08-2025 (enclose an extra copy of this form).

5. Change in Entity Status (from status indicated above)

☐ a. Applicant claims SMALL ENTITY status. See 37 CFR 1.27.☐ b. Applicant is no longer claiming SMALL ENTITY status. See 37 CFR 1.27(e)(2).

NOTE: The Issue Fee and Publication Fee (if required) will not be accepted from anyone other than the applicant; a registered attorney or agent; or the assignee or other party in interest as shown by the records of the United States Patent and Trademark Office.

Authorized Signature

LOYD E Dakin

Date

11-01-06

Typed or printed name

Registration No.

38,423

This collection of information is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PTOL-85 (Rev. 07/06) Approved for use through 04/30/2007.

OMB 0651-0033

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

PAGE 1/1 * RCVD AT 11/2/2006 11:37:46 AM [Eastern Standard Time] * SVR:USPTO-EFXXF-2/14 * DNIS:2732885 * CSID:6508528063 * DURATION (mm:ss):00:42



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

NOTICE OF ALLOWANCE AND FEE(S) DUE

22879 7590 10/24/2006

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER

MCCARTHY, CHRISTOPHER S

ART UNIT

PAPER NUMBER

2113

DATE MAILED: 10/24/2006

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	06/06/2003	Minwen Ji	200309044-1	5067

TITLE OF INVENTION: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

APPLN. TYPE	SMALL ENTITY	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	NO	\$1400	\$300	\$0	\$1700	01/24/2007

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. **PROSECUTION ON THE MERITS IS CLOSED.** THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN **THREE MONTHS** FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. **THIS STATUTORY PERIOD CANNOT BE EXTENDED.** SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE DOES NOT REFLECT A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE IN THIS APPLICATION. IF AN ISSUE FEE HAS PREVIOUSLY BEEN PAID IN THIS APPLICATION (AS SHOWN ABOVE), THE RETURN OF PART B OF THIS FORM WILL BE CONSIDERED A REQUEST TO REAPPLY THE PREVIOUSLY PAID ISSUE FEE TOWARD THE ISSUE FEE NOW DUE.

HOW TO REPLY TO THIS NOTICE:

I. Review the SMALL ENTITY status shown above.

If the SMALL ENTITY is shown as YES, verify your current SMALL ENTITY status:

A. If the status is the same, pay the TOTAL FEE(S) DUE shown above.

B. If the status above is to be removed, check box 5b on Part B - Fee(s) Transmittal and pay the PUBLICATION FEE (if required) and twice the amount of the ISSUE FEE shown above, or

If the SMALL ENTITY is shown as NO:

A. Pay TOTAL FEE(S) DUE shown above, or

B. If applicant claimed SMALL ENTITY status before, or is now claiming SMALL ENTITY status, check box 5a on Part B - Fee(s) Transmittal and pay the PUBLICATION FEE (if required) and 1/2 the ISSUE FEE shown above.

II. PART B - FEE(S) TRANSMITTAL, or its equivalent, must be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted. If an equivalent of Part B is filed, a request to reapply a previously paid issue fee must be clearly made, and delays in processing may occur due to the difficulty in recognizing the paper as an equivalent of Part B.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Utility patents issuing on applications filed on or after Dec. 12, 1980 may require payment of maintenance fees. It is patentee's responsibility to ensure timely payment of maintenance fees when due.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: **Mail** Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or **Fax** (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

22879 7590 10/24/2006

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

	(Depositor's name)
	(Signature)
	(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	06/06/2003	Minwen Ji	200309044-1	5067

TITLE OF INVENTION: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

APPLN. TYPE	SMALL ENTITY	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	NO	\$1400	\$300	\$0	\$1700	01/24/2007

EXAMINER	ART UNIT	CLASS-SUBCLASS
MCCARTHY, CHRISTOPHER S	2113	714-006000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).
☐ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.
☐ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. Use of a Customer Number is required.

2. For printing on the patent front page, list
(1) the names of up to 3 registered patent attorneys or agents OR, alternatively, 1 _____
(2) the name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed. 2 _____
3 _____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.111. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE

(B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent): ☐ Individual ☐ Corporation or other private group entity ☐ Government

4a. The following fee(s) are submitted:

- ☐ Issue Fee
☐ Publication Fee (No small entity discount permitted)
☐ Advance Order - # of Copies _____

4b. Payment of Fee(s): (Please first reapply any previously paid issue fee shown above)

- ☐ A check is enclosed.
☐ Payment by credit card. Form PTO-2038 is attached.
☐ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment, to Deposit Account Number _____ (enclose an extra copy of this form).

5. Change in Entity Status (from status indicated above)

- ☐ a. Applicant claims SMALL ENTITY status. See 37 CFR 1.27. ☐ b. Applicant is no longer claiming SMALL ENTITY status. See 37 CFR 1.27(g)(2).

NOTE: The Issue Fee and Publication Fee (if required) will not be accepted from anyone other than the applicant; a registered attorney or agent; or the assignee or other party in interest as shown by the records of the United States Patent and Trademark Office.

Authorized Signature _____

Date _____

Typed or printed name _____

Registration No. _____

This collection of information is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	06/06/2003	Minwen Ji	200309044-1	5067
22879	7590	10/24/2006	EXAMINER	
HEWLETT PACKARD COMPANY P O BOX 272400, 3404 E. HARMONY ROAD INTELLECTUAL PROPERTY ADMINISTRATION FORT COLLINS, CO 80527-2400			MCCARTHY, CHRISTOPHER S	
			ART UNIT	PAPER NUMBER
			2113	
			DATE MAILED: 10/24/2006	

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(application filed on or after May 29, 2000)

The Patent Term Adjustment to date is 581 day(s). If the issue fee is paid on the date that is three months after the mailing date of this notice and the patent issues on the Tuesday before the date that is 28 weeks (six and a half months) after the mailing date of this notice, the Patent Term Adjustment will be 581 day(s).

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) WEB site (<http://pair.uspto.gov>).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

Notice of Allowability	Application No.	Applicant(s)	
	10/456,053	JI ET AL.	
	Examiner	Art Unit	
	Christopher S. McCarthy	2113	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to 10/3/06.

2. ☒ The allowed claim(s) is/are 1, 3, 7, 10, 14, 16-21, 23, 25-28, 47, 58 renumbered as 1-18.

3. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
a) ☐ All b) ☐ Some* c) ☐ None of the:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).
* Certified copies not received: _____

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

4. ☐ A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
(a) ☐ including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached
1) ☐ hereto or 2) ☐ to Paper No./Mail Date _____.
(b) ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☐ Notice of References Cited (PTO-892)

2. ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)

3. ☐ Information Disclosure Statements (PTO/SB/08),
Paper No./Mail Date _____

4. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material

5. ☐ Notice of Informal Patent Application

6. ☐ Interview Summary (PTO-413),
Paper No./Mail Date _____

7. ☐ Examiner's Amendment/Comment

8. ☐ Examiner's Statement of Reasons for Allowance

9. ☐ Other _____


ROBERT BEAUSOLIEL
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100



PATENT
Attorney Docket No.: 200309044-1
Express Mail Label No.: EV 889716290 US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Minwen Ji <i>et al.</i>	)	Art Unit: 2113
Serial No.: 10/456,053	)	Examiner: McCarthy, Christopher S.
Filed: June 6, 2003	)	Confirmation No.: 5067
For: DATA REDUNDANCY SYSTEM AND METHOD	)	Attorney Docket No.: 200309044-1

Mail Stop AF
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

OK to
Enter
[Signature]

AMENDMENT AFTER FINAL

Sir:

In response to the Examiner's action mailed August 4, 2006, our amendments and response are as follows:


Amendments to Claims begin on page 2 of this paper.

Remarks begin on page 7 of this paper.

Express Mail Label No.: EV 889716290 US
Date of Deposit: <u>Oct. 3, 2006</u>
I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to Commissioner for Patents, P.O. Box 1450, Alexandria VA, 22313-1450.
By: <i>[Signature]</i>
Name: Michael J. Morgan

Issue Classification 	Application/Control No. 10/456,053	Applicant(s)/Patent under Reexamination JI ET AL.
	Examiner Christopher S. McCarthy	Art Unit 2113

ISSUE CLASSIFICATION													
ORIGINAL				INTERNATIONAL CLASSIFICATION									
CLASS		SUBCLASS		CLAIMED				NON-CLAIMED					
714		5		G	06	F	11	/00					/
CROSS REFERENCES								/					/
CLASS	SUBCLASS (ONE SUBCLASS PER BLOCK)							/					/
711	162							/					/
718	101							/					/
								/					/
								/					/
								/					/

Christopher McCarthy 10/18/06 (Assistant Examiner) (Date)	 ROBERT M. BENSHOOF (Primary Examiner) (Date)	Total Claims Allowed: 18
 (Legal Instruments Examiner) (Date)	O.G. Print Claim(s) 1	O.G. Print Fig. 3

☐ Claims renumbered in the same order as presented by applicant												☐ CPA		☐ T.D.		☐ R.1.47	
Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original
1	1		31		61		91		121		151		181				
	2		32		62		92		122		152		182				
2	3		33		63		93		123		153		183				
	4		34		64		94		124		154		184				
	5		35		65		95		125		155		185				
	6		36		66		96		126		156		186				
3	7		37		67		97		127		157		187				
	8		38		68		98		128		158		188				
	9		39		69		99		129		159		189				
4	10		40		70		100		130		160		190				
	11		41		71		101		131		161		191				
	12		42		72		102		132		162		192				
	13		43		73		103		133		163		193				
5	14		44		74		104		134		164		194				
	15		45		75		105		135		165		195				
6	16		46		76		106		136		166		196				
7	17	17	47		77		107		137		167		197				
8	18		48		78		108		138		168		198				
9	19		49		79		109		139		169		199				
12	20		50		80		110		140		170		200				
10	21		51		81		111		141		171		201				
	22		52		82		112		142		172		202				
11	23		53		83		113		143		173		203				
	24		54		84		114		144		174		204				
13	25		55		85		115		145		175		205				
14	26		56		86		116		146		176		206				
15	27		57		87		117		147		177		207				
16	28	18	58		88		118		148		178		208				
	29		59		89		119		149		179		209				
	30		60		90		120		150		180		210				

Christopher S. McCarthy

2113

[illegible]

INTERFERENCE SEARCHED			
Class	Subclass	Date	Examiner
PGPUB	txt srch	10/18/2006	CM

[illegible]

Interference Search

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S19 3	35	((fourth with redundan\$4) with third) with storage).clm.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:54
S19 4	0	(write with batch with overwrite).clm.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:55
S19 5	13	(write with request with overwrite).clm.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:55

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	6	((("5544347") or ("5742792") or ("5909692") or ("6092066") or ("6101497") or ("6108748"))).PN.	USPAT	OR	OFF	2006/02/09 15:51
S2	372	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:30
S14	372	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:16
S15	0	718/101.ccls. and mirrir\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:17
S16	11	718/101.ccls. and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:17
S18	85	714/6.ccls. and batch	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:03
S19	94	714/6.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:22
S20	0	("2004/0268177").URPN.	USPAT	OR	OFF	2006/02/10 11:24
S21	721	714/43.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:32
S22	23	714/43.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 14:12
S23	87	714/4.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:39
S24	13260	((batch\$3 or transaction or packet) with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S25	400	((batch\$3 or transaction or packet) with adjust\$4 with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S26	370	((batch\$3 or packet) with adjust\$4 with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S27	26	((batch\$3 or packet) with adjust\$4 with size) same (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 13:11
S28	0	("2004/0260977").URPN.	USPAT	OR	OFF	2006/02/10 12:58

EAST Search History

S29	0	("2004/0260976").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S30	0	("2004/0260972").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S31	0	("2004/0250032").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S32	0	("2004/0250031").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S33	0	("2004/0250030").URPN.	USPAT	OR	OFF	2006/02/10 13:02
S34	0	("2004/0250029").URPN.	USPAT	OR	OFF	2006/02/10 13:11
S35	10	((batch\$4) with adjust\$4 with size) same (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 13:12
S36	2	"6728484".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:04
S37	2	"6728848".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:05
S38	3	"6606694".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S39	0	os200300225760	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S40	0	us200300225760	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S41	0	"200300225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S42	0	"2003/00225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S43	0	"US2003/00225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S44	0	"US2003/0225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S45	0	"US20030225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S46	2	"20030225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S47	2	"20030135478"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:08

EAST Search History

S48	2	"20040107226"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S49	2	"20020059487"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S50	10	((("6732124") or ("5706470") or ("5140592") or ("6591351") or ("6148383"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S51	1673	714/6.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/13 14:23
S52	1332	711/162.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:41
S53	1033	711/162.ccls. and (not 714/6.ccls.)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/15 14:34
S54	373	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/16 10:59
S55	423	709/233.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:44
S56	3	"20040205382"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 10:58
S57	0	("2004/0205382").URPN.	USPAT	OR	OFF	2006/02/22 10:58
S58	0	("2004/0205382").URPN.	USPAT	OR	OFF	2006/02/22 11:12
S59	20	((("5404508") or ("5596707") or ("5758359") or ("5796934") or ("5930824") or ("6247141") or ("6785786") or ("20030084372") or ("20030225646") or ("20040078628"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 11:24
S60	16	((("5513314") or ("6266784") or ("6728898") or ("5966730") or ("6163856") or ("6182086") or ("6636988") or ("20010056554"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 11:25
S61	2	"20030147344"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 13:00
S62	4	((("20030147344") or ("20030088698"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 12:46

EAST Search History

S63	22	((("4750106") or ("5325528") or ("5504861") or ("6052797") or ("6199074") or ("6282635") or ("6327671") or ("6442709") or ("6499112") or ("6553509") or ("20030041280"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 12:47
S64	4	((("6467024") or ("5720029"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 13:49
S65	4	((("5222219") or ("6493809") or ("6502205") or ("5900020"))).PN.	USPAT	OR	OFF	2006/02/23 13:51
S66	8	((("5051887") or ("5592618") or ("6618794") or ("6671705") or ("20010008008") or ("20030145179") or ("20030182525") or ("20040128442"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 13:52
S67	24	((("6816951") or ("6694447") or ("6691245") or ("6681339") or ("5615329") or ("6549921") or ("5928367") or ("20050071708") or ("6862671") or ("6728898") or ("6052797") or ("20040193952"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 14:01
S68	12	((("6816951") or ("6694447") or ("6691245") or ("6681339") or ("5615329") or ("6549921") or ("5928367") or ("20050071708") or ("6862671") or ("6728898") or ("6052797") or ("20040193952"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 14:13
S69	3	((("20010049749") or ("20030051111") or ("6591351"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 14:14
S70	8	((("6502205") or ("6931576") or ("20030101321") or ("6405294") or ("6397307") or ("6260125") or ("5592618") or ("6496908"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 15:10
S71	0	("20330135478").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:17
S72	2	("20030135478").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:19
S73	2	("20020083120").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:19
S74	5	((("6728848") or ("6606694"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:20
S75	7	((("6391576") or ("6173377") or ("6061770"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:21

10/18/06 2:31:51 PM

Page 4

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S76	14	((("20040107226") or ("20020059487") or ("6732124") or ("5706470") or ("5140592") or ("6591351") or ("6148383")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:22
S77	10	((("20040107226") or ("20020059487") or ("6732124") or ("5706470") or ("5140592")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:25
S78	20	("4534030" "5371878" "5544347" "5555371" "5742792" "5771344" "5799141" "5889935" "5901327" "5903898" "5909548" "5933653" "6044444" "6052797" "6105029").PN. OR ("6442706").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/02/24 09:36
S79	74	"5140592"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/24 11:09
S80	2	"5140592".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/24 11:09
S81	0	"08966633".ap.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/28 15:04
S82	6	"966633".ap.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/28 15:04
S83	2	("6148383").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:48
S84	62	(failover and failback) with (memory or storage)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:53
S85	4	(failover and failback) with (memory or storage) and \$3commit\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:54
S86	4	(failover and failback) with (memory or storage) and \$3commit\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:02
S87	7	(failover) with (memory or storage) and \$3commit\$5 and failback	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:03
S88	136	714/6.ccls. and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:19
S89	15	714/6.ccls. and (failover or fail-over) and (failback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:36
S90	0	("2005/0154937").URPN.	USPAT	OR	OFF	2006/03/02 14:04

EAST Search History

S91	12	714/5.ccls. and (failover or fail-over) and (failback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:11
S92	0	("2005/0071708").URPN.	USPAT	OR	OFF	2006/03/02 14:12
S93	15	711/162.ccls. and (failover or fail-over) and (failback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:19
S94	0	718/101.ccls. and (failover or fail-over) and (failback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:19
S95	1	("2002/0099916").URPN.	USPAT	OR	OFF	2006/03/02 14:21
S96	17	("5712970" "5724501" "5761705" "5895485" "5896492" "6199105" "6247099" "6279078" "6490659" "6578158").PN. OR ("6681339").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 14:27
S97	4	((("6816951") or ("6691245") or ("5615329") or ("5928367"))).PN.	USPAT	OR	OFF	2006/03/02 14:27
S98	14	("5555371" "5592618" "6078932" "6212531" "6223304" "6292905" "6363497" "6408310" "6516394" "6549921").PN. OR ("6694447").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 14:50
S99	2	"5978565".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:30
S100	15	failover with commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S101	21	(fail-over or failover) with commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S102	23	(fail-over or failover) with \$3commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S103	30	(fail-over or failover) with \$3commit\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S104	0	("2004/0268177").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S105	0	("2004/0260977").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S106	0	("2004/0260976").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S107	0	("2004/0260972").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S108	0	("2004/0250032").URPN.	USPAT	OR	OFF	2006/03/02 15:32

10/18/06 2:31:51 PM

Page 6

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S109	0	("2004/0250030").URPN.	USPAT	OR	OFF	2006/03/02 15:33
S110	0	("2004/0250029").URPN.	USPAT	OR	OFF	2006/03/02 15:33
S111	31	714/6.ccls. and (failover or fail-over) and commit\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:57
S112	20	714/6.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:02
S113	8	714/5.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:03
S114	18	711/162.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:21
S115	7	714/6.ccls. and remirror	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:06
S116	82	714/6.ccls. and mirror\$4 and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:02
S117	146	master and slave and mirror\$4 and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:02
S118	129	(failover or fail-over) and (commit\$4 with (copy\$4 or replicat\$6 or mirror\$4))	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:07
S119	2	"6970987".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:43
S120	9	("20020010542" "20020099727" "20030105912" "5321703" "5675802" "5729733" "6334168" "6336173" "6502205").PN. OR ("6970987").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 17:43
S121	29048	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:20
S122	1603	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:20
S123	71	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory) and (\$3commit\$6)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:33

EAST Search History

S12 4	2442	(prior or original) adj2 (primary or master or first) with (storage or memory)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:28
S12 5	12	((prior or original) adj2 (primary or master or first) with (storage or memory)) same ((slave or second\$4) with (memory or storage)) and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:29
S12 6	0	("2005/0015663").URPN.	USPAT	OR	OFF	2006/03/02 19:32
S12 7	1	"5537533".pn.	USPAT	OR	OFF	2006/03/02 19:32
S12 8	22	(local) with (remote) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory) and (\$3commit\$6)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:33
S12 9	0	("2005/0198454").URPN.	USPAT	OR	OFF	2006/03/02 19:33
S13 0	75	"714"/\$3.ccls. and (failover or fail-over) and (failback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:41
S13 1	51	707/202.ccls. and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:45
S13 2	3928	(failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:46
S13 3	759	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:47
S13 4	724	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5) and \$3commit\$6 and (memory or storage)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:47
S13 5	759	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5 with (memory or storage)) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 6	755	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 7	0	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6 and mirror44	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48

EAST Search History

S13 8	343	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6 and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 9	343	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (second\$4 or remote or slave) with (memory or storage)) and \$3commit\$6 and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 13:45
S14 0	224	reintegration and storage	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 13:45
S14 1	22	reintegration and storage and (failover or fail-over or fail adj over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 13:45
S14 2	1475	714/4.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:12
S14 3	216	714/4.ccls. and (failover or fail-over or fail adj over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:34
S14 4	23	714/4.ccls. and (failback or fail-back or fail adj back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:13
S14 5	196	S143 and (not S144)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:34
S14 6	0	("2005/0229022").URPN.	USPAT	OR	OFF	2006/03/03 14:37
S14 7	0	("2005/0198552").URPN.	USPAT	OR	OFF	2006/03/03 14:38
S14 8	0	("2005/0198327").URPN.	USPAT	OR	OFF	2006/03/03 14:39
S14 9	0	("2005/0138461").URPN.	USPAT	OR	OFF	2006/03/03 14:41
S15 0	0	("2004/0158766").URPN.	USPAT	OR	OFF	2006/03/03 15:46
S15 1	441	(san or nas) and (failover or fail-over)	USPAT	OR	OFF	2006/03/03 15:46
S15 2	38	(san or nas) and (failback or fail-back)	USPAT	OR	OFF	2006/03/03 15:49
S15 3	3	"988887".ap.	USPAT	OR	OFF	2006/03/03 15:49

EAST Search History

S15 4	413	S151 and (not S152)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 16:03
S15 5	0	("6996691").URPN.	USPAT	OR	OFF	2006/03/03 17:29
S15 6	12	("5544347" "5734818" "5771344" "6065018" "6167531" "6199074" "6308284").PN. OR ("6658590").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:30
S15 7	7	("20010023471" "5274645").PN. OR ("6658540").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:39
S15 8	21	("5274645" "5544347" "5574950" "5768623" "5771344" "5790775" "5870537" "5933653" "6006342" "6167531" "6401170").PN. OR ("6643795").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:40
S15 9	12	("5274645" "5734818" "5771344" "6167531" "6199074" "6308284"). PN. OR ("6629264").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:43
S16 0	14	("5274645" "5544347" "5768623" "5790775" "6178521" "6363462" "6401170" "6457098").PN. OR ("6601187").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:44
S16 1	10	("5533190" "5721918" "5893155" "6041423" "6047294" "6148368" "6185663" "6275783" "6421688" "6658540").PN. OR ("7007042").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:49
S16 2	4	("20030041283" "20030126315" "5790775" "6574753").PN. OR ("6961867").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:51
S16 3	31	("5157663" "5566297" "5666479" "5668943" "5682534" "5737514" "5796934" "5819019" "5852724" "5907673" "5958070").PN. OR ("6185695").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:53
S16 4	54	("5077658" "5201044" "5513314" "5564011" "5649195" "5675510" "5682527" "5689706" "5765151" "5832222" "5878434" "5893140" "5909540" "5913160" "5937343" "5940826" "5948062" "5951695" "5966706" "5987506" "5987621" "5991771" "6016553").PN. OR ("6163856").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:57

EAST Search History

S16 5	18	("20010056554" "20030084372" "20030225646" "20040078628" "5404508" "5513314" "5596707" "5758359" "5796934" "5930824" "5966730" "6163856" "6182086" "6247141" "6266784" "6636988" "6728898" "6785786").PN. OR ("6990605").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:59
S16 6	2	"6173377".pn.	US-PGPUB; USPAT; DERWENT	AND	OFF	2006/07/20 12:51
S16 7	1863	714/6.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 15:57
S16 8	398	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 15:58
S16 9	4918	709/223.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 16:28
S17 0	0	(communication or message or stream or update or batch\$3) with write with expected with fail\$4 with time	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 16:29
S17 1	8	batch with transaction with overwrit\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:36
S17 2	17	batch with transaction with adjust\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:37
S17 3	423	batch with size with adjust\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:37
S17 4	10	batch with size with adjust\$4 with fail\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:37
S17 5	0	"6173377".pn. and failover	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/31 11:40
S17 6	0	"6173377".pn. and fail-over	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/31 11:40
S17 7	0	"6173377".pn. and fail adj over	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/31 11:40
S17 8	1	"6173377".pn. and switch	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 13:59

EAST Search History

S17 9	1569	711/162.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 13:59
S18 0	23	714/43.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 14:13
S18 1	774	714/43.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 14:13
S18 2	921	714/5.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/02 13:33
S18 3	6	(storage with batch with overwrit\$3). clm.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/02 13:34
S18 4	417	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:30
S18 5	3533	714/5,6,43.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:32
S18 6	1645	714/4.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:39
S18 7	1684	711/162.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:41
S18 8	479	709/233.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:45
S18 9	2361	(fourth with redundan\$4)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:46
S19 0	881	((fourth with redundan\$4) with third	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:46
S19 1	2	((fourth with redundan\$4) with third) with failover	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:47
S19 2	54	((fourth with redundan\$4) with third) with storage	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:51
S19 3	35	((fourth with redundan\$4) with third) with storage).clm.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/10/18 11:54



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	06/06/2003	Minwen Ji	200309044-1	5067

22879 7590 10/12/2006

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER

MCCARTHY, CHRISTOPHER S

ART UNIT PAPER NUMBER

2113

DATE MAILED: 10/12/2006

Please find below and/or attached an Office communication concerning this application or proceeding.



UNITED STATES DEPARTMENT OF COMMERCE
U.S. Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450

APPLICATION NO./ CONTROL NO.	FILING DATE	FIRST NAMED INVENTOR / PATENT IN REEXAMINATION	ATTORNEY DOCKET NO.
---------------------------------	-------------	---	---------------------

EXAMINER

ART UNIT	PAPER
----------	-------

20061003

DATE MAILED:

Please find below and/or attached an Office communication concerning this application or proceeding.

Commissioner for Patents

See Interview Summary.

Interview Summary	Application No.		Applicant(s)	
	10/456,053		JI ET AL.	
	Examiner		Art Unit	
	Christopher S. McCarthy		2113	

All participants (applicant, applicant's representative, PTO personnel):

(1) Christopher S. McCarthy. (3) _____

(2) Daniel Roberts, Reg. #52172. (4) _____

Date of Interview: 03 October 2006.

Type: a) ☒ Telephonic b) ☐ Video Conference
c) ☐ Personal [copy given to: 1) ☐ applicant 2) ☐ applicant's representative]

Exhibit shown or demonstration conducted: d) ☐ Yes e) ☐ No.
If Yes, brief description: _____

Claim(s) discussed: 29.

Identification of prior art discussed: Yanai et al. 6,173,377.

Agreement with respect to the claims f) ☐ was reached. g) ☒ was not reached. h) ☐ N/A.


Substance of Interview including description of the general nature of what was agreed to if an agreement was reached, or any other comments: Applicant's representative wished to discuss the difference between the batch language of the present invention compared to the write transmissions of Yanai. The examiner responded that Yanai teaches bundling the write commands into one write command before transmission and this constitutes a batch.

(A fuller description, if necessary, and a copy of the amendments which the examiner agreed would render the claims allowable, if available, must be attached. Also, where no copy of the amendments that would render the claims allowable is available, a summary thereof must be attached.)

THE FORMAL WRITTEN REPLY TO THE LAST OFFICE ACTION MUST INCLUDE THE SUBSTANCE OF THE INTERVIEW. (See MPEP Section 713.04). If a reply to the last Office action has already been filed, APPLICANT IS GIVEN A NON-EXTENDABLE PERIOD OF THE LONGER OF ONE MONTH OR THIRTY DAYS FROM THIS INTERVIEW DATE, OR THE MAILING DATE OF THIS INTERVIEW SUMMARY FORM, WHICHEVER IS LATER, TO FILE A STATEMENT OF THE SUBSTANCE OF THE INTERVIEW. See Summary of Record of Interview requirements on reverse side or on attached sheet.

Examiner Note: You must sign this form unless it is an Attachment to a signed Office action.


Examiner's signature, if required



HEWLETT-PACKARD COMPANY
Intellectual Property Administration
PO Box 272400
Fort Collins, Colorado 80527-2400

10-04-06

IFW
AF
2113

PATENT APPLICATION

ATTORNEY DOCKET NO. 200309044-1

Inventor(s): Minwen Ji et al.

Confirmation No.: 5067

Application No.: 10/456,053

Examiner: McCarthy, Christopher S.

Filing Date: June 6, 2003

Group Art Unit: 2113

Title: DATA REDUNDANCY SYSTEM AND METHOD

Mail Stop After Final
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL LETTER FOR RESPONSE/AMENDMENT

Transmitted herewith is/are the following in the above-identified application:

- ☒ Response/Amendment
☐ New fee as calculated below
☒ No additional fee
☒ Other Return postcard (2)

- ☐ Petition to extend time to respond
☐ Supplemental Declaration

Fee\$

CLAIMS AS AMENDED BY OTHER THAN A SMALL ENTITY						
(1) FOR	(2) CLAIMS REMAINING AFTER AMENDMENT	(3) NUMBER EXTRA	(4) HIGHEST NUMBER PREVIOUSLY PAID FOR	(5) PRESENT EXTRA	(6) RATE	(7) ADDITIONAL FEES
TOTAL CLAIMS	18	MINUS	47	= 0	X \$50	\$ 0
INDEP. CLAIMS	3	MINUS	3	= 0	X \$200	\$ 0
☐ FIRST PRESENTATION OF A MULTIPLE DEPENDENT CLAIM					+ \$360	\$ 0
EXTENSION FEE	☐ 1st Month \$120	☐ 2nd Month \$450	☐ 3rd Month \$1020	☐ 4th Month \$1590		\$ 0
OTHER FEES						\$
TOTAL ADDITIONAL FEE FOR THIS AMENDMENT						\$ 0

Charge \$ 0 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

Express Mail Label No.: EV 889716290 US

Date of Deposit: October 3, 2006

Typed Name: Michael J. Morgan

Signature:

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to:
Commissioner for Patents, Alexandria, VA 22313-1450

Respectfully submitted

Minwen Ji et al.

By

Daniel W. Roberts

Attorney/Agent for Applicant(s)

Reg No.: 52,172

Date: October 3, 2006

Telephone: 720-931-3016



PATENT
Attorney Docket No.: 200309044-1
Express Mail Label No.: EV 889716290 US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Minwen Ji <i>et al.</i>	)	Art Unit: 2113
Serial No.: 10/456,053	)	Examiner: McCarthy, Christopher S.
Filed: June 6, 2003	)	Confirmation No.: 5067
For: DATA REDUNDANCY SYSTEM AND METHOD	)	Attorney Docket No.: 200309044-1

Mail Stop AF
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

AMENDMENT AFTER FINAL

Sir:

In response to the Examiner's action mailed August 4, 2006, our amendments and response are as follows:


Amendments to Claims begin on page 2 of this paper.

Remarks begin on page 7 of this paper.

Express Mail Label No.: **EV 889716290 US**

Date of Deposit: Oct. 3, 2006

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to Commissioner for Patents, P.O. Box 1450, Alexandria VA, 22313-1450.

By: 

Name: Michael J. Morgan

In the Claims

Please amend the Claims of the application as indicated below.

1. (Previously Presented) A data redundancy system, comprising:
 - a primary storage facility for storing a primary copy of data, the primary storage facility including;
 - a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and
 - a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance;
 - a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility discrete from the primary storage facility and including;
 - a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and
 - a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance.
2. (Canceled)
3. (Previously Presented) The data redundancy system according to claim 1, wherein the secondary storage facility assumes the role of the primary storage facility in response to a failover event, the failover event including a communication traffic condition having a greater portion of request traffic originating closer to the secondary storage facility than to the primary storage facility.
4. – 6. (Canceled)
7. (Previously Presented) The data redundancy system according to claim 1, wherein the second redundancy appliance assumes the role of the first redundancy appliance in response to a failover event selected from the group consisting of a manually initiated event and an automatically initiated event.
8. – 9. (Canceled)

10. (Previously Presented) The data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage by forwarding a write record to the second redundancy appliance for each write request, the second redundancy appliance sending acknowledgements to the first redundancy appliance in batches in response to receiving write records.
11. – 13. (Canceled)
14. (Previously Presented) The data redundancy system according to claim 10, wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance.
15. (Canceled)
16. (Original) The data redundancy system according to claim 1, wherein the first redundancy appliance is disjoint from the second redundancy appliance.
17. (Original) The data redundancy system according to claim 1, wherein the second redundancy appliance maintains a shadow copy of data stored by the first redundancy appliance wherein the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance.
18. (Previously Presented) The data redundancy system according to claim 1, wherein the second redundancy appliance accesses the data stored in storage by the first redundancy appliance when the second redundancy appliance acts as the first redundancy appliance in response to a failover event.
19. (Previously Presented) The data redundancy system according to claim 1, wherein the first redundancy appliance forwards redundant data to the secondary storage facility as send batches.
20. (Previously Presented) The data redundancy system according to claim 1, wherein a first write transaction in a send batch at the primary storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction.
21. (Previously Presented) The data redundancy system according to claim 19, wherein a send batch size is determined according to either a time interval for collecting transactions into the send batch or a predetermined quantity of transactions to be included in the send batch.
22. (Canceled)

23. (Previously Presented) The data redundancy system according to claim 19, wherein a send batch size is determined adaptively based on available communication bandwidth between the primary storage facility and the secondary storage facility.
24. (Canceled)
25. (Previously Presented) The data redundancy system according to claim 1, wherein the third redundancy appliance stores redundant data to the storage associated with the secondary facility according to receive batches.
26. (Original) The data redundancy system according to claim 25, wherein the third redundancy appliance stores a receive batch after the receive batch is determined safe to apply.
27. (Original) The data redundancy system according to claim 26, wherein a receive batch is determined safe to apply when a data integrity test is successfully performed.
28. (Previously Presented) The data redundancy system according to claim 25, wherein a first write transaction in a first receive batch at the secondary storage facility is replaced with a second write transaction in a second receive batch wherein the second write transaction overwrites data written by the first write transaction and the first and second receive batches are combined.
29. – 46. (Canceled)

47. (Currently Amended) A data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility, the method comprising:

receiving a sequence of write requests at a first redundancy appliance at the primary storage facility;

storing data for the sequence of write requests in primary storage at the primary storage facility;

logging at least a sequence of write requests for transfer to the secondary storage facility;

forwarding the logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility batching the received write request as receive batches;

storing data associated with the receive batches in secondary storage at the secondary storage facility;

The data redundancy method according to claim 29, further comprising replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction; and

the second redundancy appliance acting as the first redundancy appliance in response to a failover event.

48. – 57. (Canceled)

58. (Currently Amended) A data redundancy system, comprising:
a primary storage facility including at least one primary redundancy appliance for
receiving a sequence of write requests and for storing data for the sequence of
write requests in storage associated with the primary storage facility, the primary
storage facility batching the write requests as send batches for transfer to a
secondary storage facility; and
the secondary storage facility disjoint from the primary storage facility including at
least one secondary redundancy appliance for receiving the sequence of write
requests and for storing data for the sequence of write requests in storage
associated with the secondary storage facility, the second storage facility
batching the received write request as receive batches, the second redundancy
appliance acting as the primary redundancy appliance in response to a failover
event;
~~The data redundancy system according to claim 51, wherein a first write transaction~~
~~in a receive batch at the second storage facility is replaced with a second write~~
~~transaction where the second write transaction overwrites data written by the first~~
~~write transaction.~~

59. – 61. (Canceled)

Remarks

Claims 1, 3, 7, 10, 14, 16-21, 23-30, 33, 34, 36, 38, 40, 41, 43-45 and 47-61 are pending in the application. Claims 1, 3, 7, 10, 14, 16-21, 23, 25-28, 47 and 58 remain after amendment. Claims 29, 30, 33, 34, 36, 38, 40, 41, 43, 44, 45, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 59, 60 and 61 are canceled without prejudice. No new matter is introduced with the amendments to the claims.

It is believed that the remarks laid out herein below attend to all rejections and further issues raised in the pending office action dated August 4, 2006.

Amendments to Place Application in Condition for Allowance

Examiner has stated that claims 1, 3, 7, 10, 14, 16-21 and 23-28 are allowed. Examiner has further stated that claims 47 and 58 are objected to as being dependent upon a rejected base claim but would be allowable if rewritten in independent form.

Applicant has therefore amended claims 47 and 58 to incorporate the limitations of the respective base claims (29 and 51) and the rejected claims have been canceled without prejudice to provide only the allowed claim remaining, thus placing the application in a condition for allowance as indicated by examiner.

Conclusion

In view of the above Remarks and the amendments to the claims, Applicant has addressed all issues raised in the Office Action dated August 4, 2006, and respectfully solicits a Notice of Allowance for claims 1, 3, 7, 10, 14, 16- 21, 23, 25-29, 47 and 58. Should any issues remain, the Examiner is encouraged to telephone the undersigned attorney.

It is believed that all of the pending claims have been addressed. However, the absence of a reply to a specific rejection, issue or comment does not signify agreement with or concession of that rejection, issue or comment. In addition, because the arguments made above may not be exhaustive, there may be reasons for patentability of any or all pending claims (or other claims) that have not been expressed. Finally, nothing in this paper should be construed as an intent to concede any issue with regard to any claim, except as specifically stated in this paper, and the amendment of any claim does not necessarily signify concession of unpatentability of the claim prior to its amendment.

Applicant believe that no further fees are due, however, should any fee be deemed necessary in connection with this Amendment and Response, the Commissioner is authorized to charge deposit account 08-2025, referencing the Attorney Docket Number 200309044-1.

Respectfully submitted,

By:



Daniel W. Roberts,
Reg. No. 52,172
LATHROP & GAGE L.C.
4845 Pearl East Circle, Suite 300
Boulder, CO 80301
Telephone: (720) 931-3016
Facsimile: (720) 931-3001

PATENT APPLICATION FEE DETERMINATION RECORD
Effective December 8, 2004

10/456053

CLAIMS AS FILED - PART I

	(Column 1)	(Column 2)
TOTAL CLAIMS		
FOR	NUMBER FILED	NUMBER EXTRA
TOTAL CHARGEABLE CLAIMS	47 minus 20 =	27
INDEPENDENT CLAIMS	2 minus 3 =	0
MULTIPLE DEPENDENT CLAIM PRESENT ☐		

* If the difference in column 1 is less than zero, enter "0" in column 2

CLAIMS AS AMENDED - PART II

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT A	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
	6-506		
Total	42 Minus	** 47	= 0
Independent	3 Minus	*** 3	= 0
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM ☐			

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT B	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
	10-306		
Total	18 Minus	** 47	= 0
Independent	3 Minus	*** 3	= 0
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM ☐			

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT C	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
Total			
Independent			
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM ☐			

SMALL ENTITY TYPE ☐ OR OTHER THAN SMALL ENTITY

RATE	FEE	RATE	FEE
BASIC FEE		BASIC FEE	750
X\$ 25=		X\$68=	486
X100=		X200=	
+180=		+360=	
TOTAL		TOTAL	\$236

SMALL ENTITY OR OTHER THAN SMALL ENTITY

RATE	ADDITIONAL FEE	RATE	ADDITIONAL FEE
X\$ 25=		X\$50=	
X100=		X200=	
+180=		+360=	
TOTAL ADDIT. FEE		TOTAL ADDIT. FEE	

RATE	ADDITIONAL FEE	RATE	ADDITIONAL FEE
X\$ 25=		X\$50=	
X100=		X200=	
+180=		+360=	
TOTAL ADDIT. FEE		TOTAL ADDIT. FEE	

RATE	ADDITIONAL FEE	RATE	ADDITIONAL FEE
X\$ 25=		X\$50=	
X100=		X200=	
+180=		+360=	



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	06/06/2003	Minwen Ji	200309044-1	5067

22879 7590 08/04/2006

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER

MCCARTHY, CHRISTOPHER S

ART UNIT PAPER NUMBER

2113

DATE MAILED: 08/04/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No. 10/456,053	Applicant(s) JI ET AL.	
	Examiner Christopher S. McCarthy	Art Unit 2113	
	-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --		

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) ☒ Responsive to communication(s) filed on 05 June 2006.

2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.

3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) ☒ Claim(s) 1,3,7,10,14,16-21,23-30,33,34,36,38,40,41,43-45 and 47-61 is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5) ☒ Claim(s) 1,3,7,10,14,16-21 and 23-28 is/are allowed.

6) ☒ Claim(s) 29,30,33,34,36,38,40,41,43-45,48-55,57 and 59-61 is/are rejected.

7) ☒ Claim(s) 47,56 and 58 is/are objected to.

8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) ☐ The specification is objected to by the Examiner.

10) ☒ The drawing(s) filed on 06 June 2003 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 a) ☐ All b) ☐ Some * c) ☐ None of:
 1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
 * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892) 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948) 3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____	4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ 5) ☐ Notice of Informal Patent Application (PTO-152) 6) ☐ Other: _____
---	--

DETAILED ACTION

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

2. Claims 29, 33, 34, 36, 38, 41, 43, 44, 48, 49-55, 57, 59, 60 are rejected under 35

U.S.C. 102(b) as being anticipated by Yanai et al. U.S. Patent 6,173,377.

As per claim 29, Yanai teaches a data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility (column 1, lines 26-30), the method comprising: receiving a sequence of write requests at a first redundancy appliance at the primary storage facility (column 5, lines 35-38); storing data for the sequence of write requests in primary storage at the primary storage facility (column 5, lines 3-5); logging at least a sequence of write requests for transfer to the secondary storage (column 5, lines 3-5); forwarding the logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility batching the received write request as receive batches (column 5, lines 35-50; column 38, lines 63-66); storing data associated with the receive batches in secondary storage at the secondary storage facility (column 2, lines 54-59); and the second

redundancy appliance acting as the first redundancy appliance in response to a failover event (column 10, lines 19-24; column 12, lines 50-64; column 31, line 65 – column 32, line 8).

As per claim 33, Yanai teaches the data redundancy method according to claim 29, further comprising the second redundancy appliance sending an acknowledgement to the first redundancy appliance in response to receiving the write requests (column 2, lines 53-59).

As per claim 34, Yanai teaches the data redundancy method according to claim 33, further comprising the first redundancy appliance sending an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance (column 3, lines 34-38).

As per claim 36, Yanai teaches the data redundancy method according to claim 29, further comprising the second redundancy appliance maintaining a shadow copy of data stored by the first redundancy appliance (column 3, lines 21-30).

As per claim 38, Yanai teaches the data redundancy method according to claim 37, further comprising determining a send batch size according to a time interval for collecting transactions into the send batch or according to a predetermined quantity of transactions to be included in the send batch (column 5, lines 39-43).

As per claim 41, Yanai teaches the data redundancy method according to claim 29, further comprising replacing a first write transaction in a send batch at the first storage facility with a second write transaction where the second write transaction overwrites data written by the first write transaction (column 42, lines 29-33; column 4, lines 9-12).

As per claim 43, Yanai teaches the data redundancy method according to claim 42, further comprising updating a shadow redundancy appliance at the secondary storage facility

Art Unit: 2113

with information about the data in storage at the secondary storage facility (column 8, lines 60-67).

As per claim 44, Yanai teaches the data redundancy method according to claim 29, further comprising storing data for the write requests in storage at the secondary storage facility according to receive batches (column 5, lines 35-50).

As per claim 48, Yanai teaches the data redundancy method according to claim 29, wherein the primary storage facility includes at least one primary redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility (column 37, lines 57-63); and wherein the secondary storage facility includes at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility (column 9, lines 33-37).

As per claim 49, Yanai teaches the data redundancy method according to claim 48, wherein the secondary storage facility includes two secondary redundancy appliances: a first secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in the storage associated with the secondary storage facility', and a second secondary redundancy appliance shadowing the first secondary redundancy appliance, the second secondary redundancy appliance acting as the first secondary redundancy appliance in response to a failover event (column 8, lines 60-67).

As per claim 50, Yanai teaches the data redundancy method according to claim 48, wherein the primary storage facility includes two primary redundancy appliances: a first primary redundancy appliance for receiving the sequence of write requests and for storing data for the

Art Unit: 2113

sequence of write requests in the storage associated with the primary storage facility, and a second primary redundancy appliance shadowing the first primary redundancy appliance, the second primary redundancy appliance acting as the first primary redundancy appliance in response to a failover event (column 8, lines 60-67; column 10, lines 19-24).

As per claim 51, Yanai teaches a data redundancy system, comprising: a primary storage facility including at least one primary redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility, the primary storage facility batching the write requests as send batches for transfer to a secondary storage facility (column 5, lines 36-50); and the secondary storage facility disjoint from the primary storage facility including at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility, the secondary storage facility batching the received write request as receive batches, the second redundancy appliance acting as the primary redundancy appliance in response to a failover event (column 5, lines 36-50; column 10, lines 19-24; column 12, lines 50-67).

As per claim 52, Yanai teaches the data redundancy system of claim 51, wherein the secondary storage facility includes two secondary redundancy appliances: a first secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in the storage associated with the secondary storage facility, and a second secondary redundancy appliance shadowing the first secondary redundancy appliance, the second secondary redundancy appliance acting as the first secondary redundancy appliance in response to a failover event (column 8, lines 60-67; column 10, lines 19-24).

As per claim 53, Yanai teaches the data redundancy system of claim 51, wherein the primary storage facility includes two primary redundancy appliances: a first primary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in the storage associated with the primary storage facility, and a second primary redundancy appliance shadowing the first primary redundancy appliance, the second primary redundancy appliance acting as the first primary redundancy appliance in response to a failover event (column 8, lines 60-67; column 10, lines 19-24).

As per claim 54, Yanai teaches the data redundancy system according to claim 51 , wherein the failover event is selected from the group consisting of a manually initiated event and an automatically initiated event (column 10, lines 15-24).

As per claim 55, Yanai teaches the data redundancy system according to claim 51, wherein the failover event is a\ communication traffic condition (column 10, lines 19-24).

As per claim 57, Yanai teaches the data redundancy system according to claim 51 , wherein a first write transaction in a send batch at the first storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction (column 42, lines 29-33; column 4, lines 9-12).

As per claim 59, Yanai teaches the data redundancy system according to claim 51 , wherein the primary redundancy appliance updates the secondary redundancy appliance with information about the data in storage associated with the primary storage facility by forwarding a write record to the secondary redundancy appliance for each write request, the secondary redundancy appliance sending acknowledgements to the primary redundancy appliance in batches in response to receiving write records (column 41, lines 16-53).

As per claim 60, Yanai teaches the data redundancy system according to claim 51, wherein the secondary redundancy appliance accesses data stored in the storage associated with the primary storage facility by the primary redundancy appliance when the secondary redundancy appliance acts as the primary redundancy appliance in response to the failover event (column 31, line 65 – column, 32, line 8).

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. Claim 30 is rejected under 35 U.S.C. 103(a) as being unpatentable over Yanai in view of Leach et al. U.S. Patent 6,694,447.

As per claim 30, Yanai teaches the data redundancy method according to claim 29. Yanai does not explicitly teach wherein said first redundancy appliance performs said primary storing prior to a failover event and wherein the second redundancy appliance performs primary storing after the failover event.(column 6, lines 19-35). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the storing process of Leach in the redundancy process of Yanai. One of ordinary skill would have been motivated to use the storing process of Leach in the redundancy process of Yanai because Leach teaches the importance of remote storage recovery used for disaster protection (column 1, lines 24-25); an explicit desire of Yanai (column 1, lines 38-46; column 2, lines 31-52).

5. Claims 40, 45, 61 are rejected under 35 U.S.C. 103(a) as being unpatentable over Yanai in view of Wahl et al. U.S. Patent 6,442,706.

As per claim 40, Yanai teaches the data redundancy method according to claim 29. Yanai does not explicitly teach it further comprising determining a send batch size adaptively based on available communication bandwidth between the primary storage facility and the secondary storage facility. Wahl does teach determining a send batch size adaptively based on available communication bandwidth between the primary storage facility and the secondary storage facility (14, lines 53-55). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the communication process of Wahl in the recovery process of Yanai. One of ordinary skill in the art would have been motivated to use the communication process of Wahl in the recovery process of Yanai because Wahl teaches the importance of the restoration of data in disaster recovery (column 1, lines 12-15); an explicit desire of Yanai (column 1, lines 38-46; column 2, lines 31-52).

As per claim 45, Yanai teaches the data redundancy method according to claim 44. Yanai does not teach it further comprising determining whether a receive batch is safe to apply before storing the receive batch by performing a data integrity test on data of the receive batch. Wahl does teach determining whether a receive batch is safe to apply before storing the receive batch by performing a data integrity test on data of the receive batch (column 24, lines 34-42). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the communication process of Wahl in the recovery process of Yanai. One of ordinary skill in the art would have been motivated to use the communication process of Wahl in the

Art Unit: 2113

recovery process of Yanai because Wahl teaches the importance of the restoration of data in disaster recovery (column 1, lines 12-15); an explicit desire of Yanai (column 1, lines 38-46; column 2, lines 31-52).

As per claim 61, Yanai teaches the data redundancy system according to claim 51.

Yanai does not teach wherein a send batch size is adaptively determined based on available communication bandwidth between the primary storage facility and the secondary storage facility. Wahl teaches wherein a send batch size is adaptively determined based on available communication bandwidth between the primary storage facility and the secondary storage facility (column 14, lines 53-55). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the communication process of Wahl in the recovery process of Yanai. One of ordinary skill in the art would have been motivated to use the communication process of Wahl in the recovery process of Yanai because Wahl teaches the importance of the restoration of data in disaster recovery (column 1, lines 12-15); an explicit desire of Yanai (column 1, lines 38-46; column 2, lines 31-52).

Allowable Subject Matter

6. Claims 47, 56, 58 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

7. Claims 1, 3, 7, 10, 14, 16-21, 23-28 are allowed.

Conclusion

8. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Christopher S. McCarthy whose telephone number is (571)272-3651. The examiner can normally be reached on M-F, 9 - 5:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Robert Beausoliel can be reached on (571)272-3645. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2113

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

csm


ROBERT M. BEAUCHAMP
SUPERVISORY PATENT EXAMINER
U.S. PATENT AND TRADEMARK OFFICE

Notice of References Cited	Application/Control No. 10/456,053	Applicant(s)/Patent Under Reexamination JI ET AL.	
	Examiner Christopher S. McCarthy	Art Unit 2113	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-6,173,377 B1	01-2001	Yanai et al.	711/162
	B	US-			
	C	US-			
	D	US-			
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			


FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Index of Claims 				Application/Control No. 10/456,053 Examiner Christopher S. McCarthy		Applicant(s)/Patent under Reexamination JI ET AL. Art Unit 2113	
---	--	--	--	--	--	--	--

√ Rejected = Allowed	- (Through numeral) Cancelled + Restricted	N Non-Elected I Interference	A Appeal O Objected
-------------------------	---	---------------------------------	------------------------

Claim		Date						
Final	Original	8/1/06						
1	=							
2								
3	=							
4								
5								
6								
7	=							
8								
9								
10	=							
11								
12								
13								
14	=							
15								
16	=							
17	=							
18	=							
19	=							
20	=							
21	=							
22								
23	=							
24	=							
25	=							
26	=							
27	=							
28	=							
29	√							
30	√							
31								
32								
33	√							
34	√							
35								
36	√							
37								
38	√							
39								
40	√							
41	√							
42								
43	√							
44	√							
45	√							
46								
47	o							
48	√							
49	√							
50	√							

Claim		Date						
Final	Original	8/1/06						
51	√							
52	√							
53	√							
54	√							
55	√							
56	o							
57	√							
58	o							
59	√							
60	√							
61	√							
62								
63								
64								
65								
66								
67								
68								
69								
70								
71								
72								
73								
74								
75								
76								
77								
78								
79								
80								
81								
82								
83								
84								
85								
86								
87								
88								
89								
90								
91								
92								
93								
94								
95								
96								
97								
98								
99								
100								

Claim		Date						
Final	Original							
101								
102								
103								
104								
105								
106								
107								
108								
109								
110								
111								
112								
113								
114								
115								
116								
117								
118								
119								
120								
121								
122								
123								
124								
125								
126								
127								
128								
129								
130								
131								
132								
133								
134								
135								
136								
137								
138								
139								
140								
141								
142								
143								
144								
145								
146								
147								
148								
149								
150								

Search Notes 	Application/Control No. 10/456,053 Examiner Christopher S. McCarthy	Applicant(s)/Patent under Reexamination JI ET AL. Art Unit 2113
--	--	--

SEARCHED			
Class	Subclass	Date	Examiner
711	162	8/1/2006	CM
709	233	7/28/2006	CM
714	6	7/28/2006	CM
718	101	7/28/2006	CM

INTERFERENCE SEARCHED			
Class	Subclass	Date	Examiner

SEARCH NOTES (INCLUDING SEARCH STRATEGY)		
	DATE	EXMR
714/5,43 - text search only - see attached search report	8/1/2006	CM
Search updated	8/1/2006	CM

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	6	((("5544347") or ("5742792") or ("5909692") or ("6092066") or ("6101497") or ("6108748")).PN.	USPAT	OR	OFF	2006/02/09 15:51
S2	372	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:21
S14	372	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:16
S15	0	718/101.ccls. and mirrir\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:17
S16	11	718/101.ccls. and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:17
S18	85	714/6.ccls. and batch	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:03
S19	94	714/6.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:22
S20	0	("2004/0268177").URPN.	USPAT	OR	OFF	2006/02/10 11:24
S21	721	714/43.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 12:33
S22	23	714/43.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 14:12
S23	87	714/4.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 12:56
S24	13260	((batch\$3 or transaction or packet) with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S25	400	((batch\$3 or transaction or packet) with adjust\$4 with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S26	370	((batch\$3 or packet) with adjust\$4 with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S27	26	((batch\$3 or packet) with adjust\$4 with size) same (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 13:11
S28	0	("2004/0260977").URPN.	USPAT	OR	OFF	2006/02/10 12:58

EAST Search History

S29	0	("2004/0260976").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S30	0	("2004/0260972").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S31	0	("2004/0250032").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S32	0	("2004/0250031").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S33	0	("2004/0250030").URPN.	USPAT	OR	OFF	2006/02/10 13:02
S34	0	("2004/0250029").URPN.	USPAT	OR	OFF	2006/02/10 13:11
S35	10	((batch\$4) with adjust\$4 with size) same (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 13:12
S36	2	"6728484".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:04
S37	2	"6728848".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:05
S38	3	"6606694".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S39	0	os200300225760	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S40	0	us200300225760	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S41	0	"200300225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S42	0	"2003/00225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S43	0	"US2003/00225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S44	0	"US2003/0225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S45	0	"US20030225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S46	2	"20030225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S47	2	"20030135478"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:08

EAST Search History

S48	2	"20040107226"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S49	2	"20020059487"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S50	10	((("6732124") or ("5706470") or ("5140592") or ("6591351") or ("6148383"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S51	1673	714/6.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/13 14:23
S52	1332	711/162.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/14 17:28
S53	1033	711/162.ccls. and (not 714/6.ccls.)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/15 14:34
S54	373	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/16 10:59
S55	423	709/233.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/16 11:00
S56	3	"20040205382"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 10:58
S57	0	("2004/0205382").URPN.	USPAT	OR	OFF	2006/02/22 10:58
S58	0	("2004/0205382").URPN.	USPAT	OR	OFF	2006/02/22 11:12
S59	20	((("5404508") or ("5596707") or ("5758359") or ("5796934") or ("5930824") or ("6247141") or ("6785786") or ("20030084372") or ("20030225646") or ("20040078628"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 11:24
S60	16	((("5513314") or ("6266784") or ("6728898") or ("5966730") or ("6163856") or ("6182086") or ("6636988") or ("20010056554"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 11:25
S61	2	"20030147344"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 13:00
S62	4	((("20030147344") or ("20030088698"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 12:46

EAST Search History

S63	22	((("4750106") or ("5325528") or ("5504861") or ("6052797") or ("6199074") or ("6282635") or ("6327671") or ("6442709") or ("6499112") or ("6553509") or ("20030041280"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 12:47
S64	4	((("6467024") or ("5720029"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 13:49
S65	4	((("5222219") or ("6493809") or ("6502205") or ("5900020"))).PN.	USPAT	OR	OFF	2006/02/23 13:51
S66	8	((("5051887") or ("5592618") or ("6618794") or ("6671705") or ("20010008008") or ("20030145179") or ("20030182525") or ("20040128442"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 13:52
S67	24	((("6816951") or ("6694447") or ("6691245") or ("6681339") or ("5615329") or ("6549921") or ("5928367") or ("20050071708") or ("6862671") or ("6728898") or ("6052797") or ("20040193952"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 14:01
S68	12	((("6816951") or ("6694447") or ("6691245") or ("6681339") or ("5615329") or ("6549921") or ("5928367") or ("20050071708") or ("6862671") or ("6728898") or ("6052797") or ("20040193952"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 14:13
S69	3	((("20010049749") or ("20030051111") or ("6591351"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 14:14
S70	8	((("6502205") or ("6931576") or ("20030101321") or ("6405294") or ("6397307") or ("6260125") or ("5592618") or ("6496908"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 15:10
S71	0	("20330135478").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:17
S72	2	("20030135478").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:19
S73	2	("20020083120").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:19

EAST Search History

S74	5	((("6728848") or ("6606694")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:20
S75	7	((("6391576") or ("6173377") or ("6061770")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:21
S76	14	((("20040107226") or ("20020059487") or ("6732124") or ("5706470") or ("5140592") or ("6591351") or ("6148383")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:22
S77	10	((("20040107226") or ("20020059487") or ("6732124") or ("5706470") or ("5140592")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:25
S78	20	("4534030" "5371878" "5544347" "5555371" "5742792" "5771344" "5799141" "5889935" "5901327" "5903898" "5909548" "5933653" "6044444" "6052797" "6105029").PN. OR ("6442706").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/02/24 09:36
S79	74	"5140592"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/24 11:09
S80	2	"5140592".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/24 11:09
S81	0	"08966633".ap.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/28 15:04
S82	6	"966633".ap.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/28 15:04
S83	2	("6148383").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:48
S84	62	(failover and failback) with (memory or storage)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:53
S85	4	(failover and failback) with (memory or storage) and \$3commit\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:54
S86	4	(failover and failback) with (memory or storage) and \$3commit\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:02
S87	7	(failover) with (memory or storage) and \$3commit\$5 and failback	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:03

EAST Search History

S88	136	714/6.ccls. and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:19
S89	15	714/6.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:36
S90	0	("2005/0154937").URPN.	USPAT	OR	OFF	2006/03/02 14:04
S91	12	714/5.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:11
S92	0	("2005/0071708").URPN.	USPAT	OR	OFF	2006/03/02 14:12
S93	15	711/162.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:19
S94	0	718/101.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:19
S95	1	("2002/0099916").URPN.	USPAT	OR	OFF	2006/03/02 14:21
S96	17	("5712970" "5724501" "5761705" "5895485" "5896492" "6199105" "6247099" "6279078" "6490659" "6578158").PN. OR ("6681339").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 14:27
S97	4	((("6816951") or ("6691245") or ("5615329") or ("5928367"))).PN.	USPAT	OR	OFF	2006/03/02 14:27
S98	14	("5555371" "5592618" "6078932" "6212531" "6223304" "6292905" "6363497" "6408310" "6516394" "6549921").PN. OR ("6694447").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 14:50
S99	2	"5978565".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:30
S10 0	15	failover with commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 1	21	(fail-over or failover) with commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 2	23	(fail-over or failover) with \$3commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 3	30	(fail-over or failover) with \$3commit\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 4	0	("2004/0268177").URPN.	USPAT	OR	OFF	2006/03/02 15:32

8/1/06 4:40:46 PM

Page 6

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S10 5	0	("2004/0260977").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 6	0	("2004/0260976").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 7	0	("2004/0260972").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 8	0	("2004/0250032").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 9	0	("2004/0250030").URPN.	USPAT	OR	OFF	2006/03/02 15:33
S11 0	0	("2004/0250029").URPN.	USPAT	OR	OFF	2006/03/02 15:33
S11 1	31	714/6.ccls. and (failover or fail-over) and commit\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:57
S11 2	20	714/6.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:02
S11 3	8	714/5.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:03
S11 4	18	711/162.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:21
S11 5	7	714/6.ccls. and remirror	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:06
S11 6	82	714/6.ccls. and mirror\$4 and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:02
S11 7	146	master and slave and mirror\$4 and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:02
S11 8	129	(failover or fail-over) and (commit\$4 with (copy\$4 or replicat\$6 or mirror\$4))	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:07
S11 9	2	"6970987".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:43
S12 0	9	("20020010542" "20020099727" "20030105912" "5321703" "5675802" "5729733" "6334168" "6336173" "6502205").PN. OR ("6970987").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 17:43
S12 1	29048	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:20

8/1/06 4:40:46 PM

Page 7

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S12 2	1603	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:20
S12 3	71	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory) and (\$3commit\$6)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:33
S12 4	2442	(prior or original) adj2 (primary or master or first) with (storage or memory)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:28
S12 5	12	((prior or original) adj2 (primary or master or first) with (storage or memory)) same ((slave or second\$4) with (memory or storage)) and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:29
S12 6	0	("2005/0015663").URPN.	USPAT	OR	OFF	2006/03/02 19:32
S12 7	1	"5537533".pn.	USPAT	OR	OFF	2006/03/02 19:32
S12 8	22	(local) with (remote) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory) and (\$3commit\$6)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:33
S12 9	0	("2005/0198454").URPN.	USPAT	OR	OFF	2006/03/02 19:33
S13 0	75	"714"/\$3.ccls. and (failover or fail-over) and (failback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:41
S13 1	51	707/202.ccls. and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:45
S13 2	3928	(failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:46
S13 3	759	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:47
S13 4	724	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5) and \$3commit\$6 and (memory or storage)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:47
S13 5	759	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5 with (memory or storage)) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48

EAST Search History

S13 6	755	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 7	0	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6 and mirror44	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 8	343	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6 and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 9	343	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (second\$4 or remote or slave) with (memory or storage)) and \$3commit\$6 and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 13:45
S14 0	224	reintegration and storage	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 13:45
S14 1	22	reintegration and storage and (failover or fail-over or fail adj over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 13:45
S14 2	1475	714/4.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:12
S14 3	216	714/4.ccls. and (failover or fail-over or fail adj over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:34
S14 4	23	714/4.ccls. and (failback or fail-back or fail adj back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:13
S14 5	196	S143 and (not S144)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 14:34
S14 6	0	("2005/0229022").URPN.	USPAT	OR	OFF	2006/03/03 14:37
S14 7	0	("2005/0198552").URPN.	USPAT	OR	OFF	2006/03/03 14:38
S14 8	0	("2005/0198327").URPN.	USPAT	OR	OFF	2006/03/03 14:39
S14 9	0	("2005/0138461").URPN.	USPAT	OR	OFF	2006/03/03 14:41
S15 0	0	("2004/0158766").URPN.	USPAT	OR	OFF	2006/03/03 15:46

8/1/06 4:40:46 PM

Page 9

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S15 1	441	(san or nas) and (failover or fail-over)	USPAT	OR	OFF	2006/03/03 15:46
S15 2	38	(san or nas) and (failback or fail-back)	USPAT	OR	OFF	2006/03/03 15:49
S15 3	3	"988887".ap.	USPAT	OR	OFF	2006/03/03 15:49
S15 4	413	S151 and (not S152)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/03 16:03
S15 5	0	("6996691").URPN.	USPAT	OR	OFF	2006/03/03 17:29
S15 6	12	("5544347" "5734818" "5771344" "6065018" "6167531" "6199074" "6308284").PN. OR ("6658590").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:30
S15 7	7	("20010023471" "5274645").PN. OR ("6658540").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:39
S15 8	21	("5274645" "5544347" "5574950" "5768623" "5771344" "5790775" "5870537" "5933653" "6006342" "6167531" "6401170").PN. OR ("6643795").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:40
S15 9	12	("5274645" "5734818" "5771344" "6167531" "6199074" "6308284").PN. OR ("6629264").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:43
S16 0	14	("5274645" "5544347" "5768623" "5790775" "6178521" "6363462" "6401170" "6457098").PN. OR ("6601187").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:44
S16 1	10	("5533190" "5721918" "5893155" "6041423" "6047294" "6148368" "6185663" "6275783" "6421688" "6658540").PN. OR ("7007042").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:49
S16 2	4	("20030041283" "20030126315" "5790775" "6574753").PN. OR ("6961867").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:51
S16 3	31	("5157663" "5566297" "5666479" "5668943" "5682534" "5737514" "5796934" "5819019" "5852724" "5907673" "5958070").PN. OR ("6185695").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:53

EAST Search History

S16 4	54	("5077658" "5201044" "5513314" "5564011" "5649195" "5675510" "5682527" "5689706" "5765151" "5832222" "5878434" "5893140" "5909540" "5913160" "5937343" "5940826" "5948062" "5951695" "5966706" "5987506" "5987621" "5991771" "6016553").PN. OR ("6163856"). URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:57
S16 5	18	("20010056554" "20030084372" "20030225646" "20040078628" "5404508" "5513314" "5596707" "5758359" "5796934" "5930824" "5966730" "6163856" "6182086" "6247141" "6266784" "6636988" "6728898" "6785786").PN. OR ("6990605"). URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/03 17:59
S16 6	2	"6173377".pn.	US-PGPUB; USPAT; DERWENT	AND	OFF	2006/07/20 12:51
S16 7	1863	714/6.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 15:57
S16 8	398	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 15:58
S16 9	4918	709/223.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 16:28
S17 0	0	(communication or message or stream or update or batch\$3) with write with expected with fail\$4 with time	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/28 16:29
S17 1	8	batch with transaction with overwrit\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:36
S17 2	17	batch with transaction with adjust\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:37
S17 3	423	batch with size with adjust\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:37
S17 4	10	batch with size with adjust\$4 with fail\$4	US-PGPUB; USPAT; DERWENT	OR	ON	2006/07/28 16:37

EAST Search History

S17 5	0	"6173377".pn. and failover	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/31 11:40
S17 6	0	"6173377".pn. and fail-over	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/31 11:40
S17 7	0	"6173377".pn. and fail adj over	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/07/31 11:40
S17 8	1	"6173377".pn. and switch	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 13:59
S17 9	1569	711/162.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 13:59
S18 0	23	714/43.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 14:13
S18 1	774	714/43.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 14:13
S18 2	921	714/5.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/08/01 14:13

06-06-06

TFW 2/13

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400



Inventor(s): Minwen Ji et al.

Application No.: 10/456,053

Filing Date: June 6, 2003

PATENT APPLICATION

ATTORNEY DOCKET NO. 200309044-1

Confirmation No.: 5067

Examiner: McCarthy, Christopher S.

Group Art Unit: 2113

Title: Fault-Tolerant Data Redundancy Technique

Mail Stop Amendment
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL LETTER FOR RESPONSE/AMENDMENT

Transmitted herewith is/are the following in the above-identified application:

- ☒ Response/Amendment
☐ New fee as calculated below
☐ No additional fee
☒ Other Return postcard (2)

- ☐ Petition to extend time to respond
☐ Supplemental Declaration

Fee\$

CLAIMS AS AMENDED BY OTHER THAN A SMALL ENTITY						
(1) FOR	(2) CLAIMS REMAINING AFTER AMENDMENT	(3) NUMBER EXTRA	(4) HIGHEST NUMBER PREVIOUSLY PAID FOR	(5) PRESENT EXTRA	(6) RATE	(7) ADDITIONAL FEES
TOTAL CLAIMS	42	MINUS	47	= 0	X \$50	\$ 0
INDEP. CLAIMS	3	MINUS	3	= 0	X \$200	\$ 0
☐ FIRST PRESENTATION OF A MULTIPLE DEPENDENT CLAIM					+ \$360	\$ 0
EXTENSION FEE	☐ 1st Month \$120	☐ 2nd Month \$450	☐ 3rd Month \$1020	☐ 4th Month \$1590		\$ 0
OTHER FEES						\$
TOTAL ADDITIONAL FEE FOR THIS AMENDMENT						\$ 0

Charge \$ 0 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

Express Mail Label No.: EV 498084825 US

Date of Deposit: June 5, 2006

Typed Name: MiMi Nguyen

Signature: Minwen Ji et al.

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to:
Commissioner for Patents, Alexandria, VA 22313-1450

Respectfully submitted,

Minwen Ji et al.

By

Daniel W. Roberts

Daniel W. Roberts

Attorney/Agent for Applicant(s)

Reg No.: 52,172

Date: June 5, 2006

Telephone: (720) 931-3016



PATENT
Attorney Docket No.: 200309044-1
Express Mail Label No.: EV 498084825 US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Minwen Ji <i>et al.</i>	)	Art Unit: 2113
Serial No.: 10/456,053	)	Examiner: McCarthy, Christopher S.
Filed: June 6, 2003	)	Confirmation No.: 5067
For: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE	)	Attorney Docket No.: 200309044-1

Mail Stop Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

AMENDMENT AND RESPONSE


Sir:

In response to the Examiner's action mailed March 10, 2006, our amendments and response are as follows:

Amendments to the Specification begin on page 2 of this paper.

Amendments to Claims begin on page 4 of this paper.

Remarks begin on page 11 of this paper.

Express Mail Label No.: EV 498084825 US
Date of Deposit : <u>June 05, 2006</u>
I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to Commissioner for Patents, P.O. Box 1450, Alexandria VA, 22313-1450.
By: 
Name: MiMi Nguyen

In the Specification

Please amend the title of the application appearing on page 1 as follows:

~~FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE~~
DATA REDUNDANCY SYSTEM AND METHOD

In accordance with 37 C.F.R. § 1.121(b)(1), Applicant presents the following replacement paragraphs to the specification:

Please amend the paragraph located on page 23, lines 8-13 as follows:

During normal operation in which the primary facility 102 is in state 402 and the secondary facility 104 is in state 404, a fault may occur in which the secondary facility 104 becomes inoperative. Under these circumstances, the secondary facility 104 may enter the failed state 406. Upon recovery, the secondary facility 104 may return to the normal state 404 directly, or via the pending state ~~412-424~~ or the update state ~~414-430~~, as described above.

Please amend the paragraph located on pages 27-28, lines 25-3 as follows:

The backup 508 is preferably sent to the local facility as a single batch, which is committed at the local facility as a whole. As explained above in reference to Figure 4-~~7~~, propagation of the backup 508 may be performed while the local facility is in the pending state 424. Note that versions A1, B1 and C1 are omitted from the backup 508 since they are not present at the remote facility. Thus, at the time T_2 , the data at both facilities is consistent, while some data has been lost (i.e. A1, B1 and C1). In particular, at time T_2 each facility has committed data 510, including versions A0, C0 and B2. Because the operations that result in data consistency may be performed under control of the state machines 400 (Figure 7), they may be performed automatically. For example, forwarding the update 508 may be performed in pending state 424 of Figure 7, while the local facility is preparing to resume operation as the secondary. Thus, operation from a state of data consistency can proceed without requiring human intervention.

Please amend the paragraph located on page 30, lines 20-29 as follows:

Though redundant logs may be maintained at each redundancy appliance ~~the-in~~ primary storage 102, both of the primary redundancy appliances 202 and 202' preferably have access to a common log 206. This is shown in Figure 11. The common log 206 may be accessible from both appliances 202 and 202', for example, via a storage area network (SAN). Maintaining a single log 206 has an advantage of avoiding any latency associated with making two synchronous updates to a shadow log. Rather, maintaining the shadow appliance's state only incurs extra latency of transferring update messages to the shadowing redundancy appliance (e.g., across an interconnecting LAN between the appliances 202 and

202'). The log 206 can be made fault-tolerant by using a local redundancy scheme, such as RAID5 or mirroring.

Please amend the paragraph located on page 33, lines 18-22 as follows:

So that the secondary facility ~~402-104~~ is able to accommodate the same workload as the primary facility 102, the secondary facility 104 is preferably also provided with additional components corresponding to those of the primary facility 102. As shown in Figure 12, the secondary facility 104 includes the appliance 204 that uses log 210 and LU 212 and an additional appliance 204" that uses log 210" and LU 212".

Please amend the paragraph located on page 35, lines 1-6 as follows:

Once the batches are terminated, each appliance 202 and 202" may then send its batch to a corresponding one of the appliances at the secondary facility ~~402-104~~. More particularly, data stored in LU 208 by appliance 202 at the primary 102 is forwarded to the appliance 204 for storage in the LU 212 at the secondary 104. In addition, data stored in the LU 208" by appliance 202" at the primary 102 is forwarded to the appliance 204" for storage in the LU 212" at the secondary 104.

Please amend the paragraph located on page 35, lines 13-29 as follows:

In addition, to maintain data consistency at the LUs of the secondary facility 204, the appliances 204 and 204" preferably coordinate committing of the batches to their respective LUs 212 and 212". Committing the batches at the secondary may also be performed as a two-phase operation. For example, one of the appliances 204 or 204" may first determine a time at which to commit the current batch according to any of the criteria described herein (e.g., after a data check). The appliance at the secondary 104 which determined the time to commit the batch may then declare this to the other appliance at the secondary ~~402-104~~ by informing the other of the time at which its batch should be committed. In this way, each appliance 204 and 204" commits its batches to storage at approximately the same point in time. Well-known techniques such as "two-phase commit" can be used to ensure that both commit, or neither does. An acknowledgement may then be sent to the primary facility 102 indicating that the data was successfully stored by the secondary facility 104. The appliances 202 and 202" at the primary 102 may then delete the corresponding records in their logs 206 and 206". Deleting of the logs at the primary 102 can be performed independently (i.e. the appliances 202 and 202" need not coordinate deletion with each other).

In the Claims

Please amend the Claims of the application as indicated below.

1. (Currently Amended) A data redundancy system, ~~including a primary facility for storing a primary copy of data and a secondary facility for storing data that is redundant of the primary copy of the data, wherein the primary facility comprises:~~ comprising:
a primary storage facility for storing a primary copy of data, the primary storage facility including:
a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and
a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance;
a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility discrete from the primary storage facility and including:
a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and
a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance.
2. (Canceled)
3. (Currently Amended) The data redundancy system according to claim 1, wherein the ~~second redundancy appliance~~ secondary storage facility assumes the role of the first redundancy appliance ~~primary storage facility~~ in response to a failover event, the failover event including a communication traffic condition having a greater portion of request traffic originating closer to the secondary storage facility than to the primary storage facility.
4. (Canceled)
5. (Canceled)
6. (Canceled)

7. (Currently Amended) The data redundancy system according to claim [6]1, wherein the second redundancy appliance assumes the role of the first redundancy appliance in response to a failover event wherein the failover event is selected from the group consisting of a manually initiated event and an automatically initiated event.
8. (Canceled)
9. (Canceled)
10. (Currently Amended) The data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage by forwarding a write record to the second redundancy appliance for each write request, the second redundancy appliance sending acknowledgements to the first redundancy appliance in batches in response to receiving write records.
11. (Canceled)
12. (Canceled)
13. (Canceled)
14. (Currently Amended) The data redundancy system according to claim [13]10, wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance.
15. (Canceled)
16. (Original) The data redundancy system according to claim 1, wherein the first redundancy appliance is disjoint from the second redundancy appliance.
17. (Original) The data redundancy system according to claim 1, wherein the second redundancy appliance maintains a shadow copy of data stored by the first redundancy appliance wherein the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance.
18. (Currently Amended) The data redundancy system according to claim 1, wherein the second redundancy appliance accesses the data stored in storage by the first redundancy appliance when the second redundancy appliance ~~assumes the role of~~ acts as the first redundancy appliance in response to a failover event.
19. (Currently Amended) The data redundancy system according to claim 1, wherein the first redundancy appliance forwards redundant data to the secondary storage facility ~~according to~~ as send batches.

20. (Currently Amended) The data redundancy system according to claim 1, wherein a first write transaction in a send batch at the ~~first-primary~~ storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction.
21. (Currently Amended) The data redundancy system according to claim [20]19, wherein a send batch size is determined according to either a time interval for collecting transactions into the send batch or a predetermined quantity of transactions to be included in the send batch.
22. (Canceled)
23. (Currently Amended) The data redundancy system according to claim [20]19, wherein a send batch size is determined adaptively based on available communication bandwidth between the ~~first-primary~~ storage facility and the ~~second-secondary~~ storage facility.
24. (Canceled)
25. (Currently Amended) The data redundancy system according to claim [24]1, wherein the third redundancy appliance stores redundant data to the storage associated with the secondary facility according to receive batches.
26. (Original) The data redundancy system according to claim 25, wherein the third redundancy appliance stores a receive batch after the receive batch is determined safe to apply.
27. (Original) The data redundancy system according to claim 26, wherein a receive batch is determined safe to apply when a data integrity test is successfully performed.
28. (Currently Amended) The data redundancy system according to claim 25, wherein a first write transaction in a first receive batch at the ~~second-secondary~~ storage facility is replaced with a second write transaction in a second receive batch wherein the second write transaction overwrites data written by the first write transaction and the first and second receive batches are combined.

29. (Currently Amended) A data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility, the method comprising:
- receiving a sequence of write requests at a first redundancy appliance at the primary storage facility;
 - storing data for the sequence of write requests in primary storage at the primary storage facility;
 - logging at least a sequence of write requests for transfer to the secondary storage facility;
 - ~~updating a shadow redundancy appliance at the primary storage facility with information about the data in storage; and~~
 - ~~forwarding the sequence of logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility batching the received write request as receive batches;~~
 - storing data associated with the receive batches in secondary storage at the secondary storage facility; and
 - the second redundancy appliance acting as the first redundancy appliance in response to a failover event.
30. (Currently Amended) The data redundancy method according to claim 29, wherein said first redundancy appliance performs ~~said primary storing and arranging~~ prior to a failover event and wherein the second redundancy appliance performs ~~said primary storing and arranging~~ after the failover event.
31. (Canceled)
32. (Canceled)
33. (Currently Amended) The data redundancy method according to claim [32]29, further comprising the second redundancy appliance sending an acknowledgement to the first redundancy appliance in response to receiving the write ~~record~~requests.
34. (Original) The data redundancy method according to claim 33, further comprising the first redundancy appliance sending an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance.
35. (Canceled)
36. (Original) The data redundancy method according to claim 29, further comprising the second redundancy appliance maintaining a shadow copy of data stored by the first redundancy appliance.
37. (Canceled)

38. (Currently Amended) The data redundancy method according to claim [37]~~29~~, further comprising determining a send batch size according to either a time interval for collecting transactions into the send batch or according to a predetermined quantity of transactions to be included in the send batch.
39. (Canceled)
40. (Currently Amended) The data redundancy method according to claim [37]~~29~~, further comprising determining a send batch size adaptively based on available communication bandwidth between the ~~first~~-primary storage facility and the ~~second~~-secondary storage facility.
41. (Currently Amended) The data redundancy method according to claim 29, further comprising replacing a first write transaction in a send batch at the ~~first~~-primary storage facility with a second write transaction where the second write transaction overwrites data written by the first write transaction.
42. (Canceled)
43. (Currently Amended) The data redundancy method according to claim [42]~~29~~, further comprising updating a shadow redundancy appliance at the secondary storage facility with information about the data in storage at the secondary storage facility.
44. (Currently Amended) The data redundancy method according to claim 29, further comprising storing data for the write requests in secondary storage at the secondary storage facility according to receive batches.
45. (Currently Amended) The data redundancy method according to claim 44, further comprising determining whether a receive batch is safe to apply before storing the receive batch by performing a data integrity test on data of the receive batch.
46. (Canceled)
47. (Currently Amended)The data redundancy method according to claim [46]~~29~~, further comprising replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction.
48. (New) The data redundancy method according to claim 29, wherein the primary storage facility includes at least one primary redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and wherein the secondary storage facility includes at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility.

49. (New) The data redundancy method according to claim 48, wherein the secondary storage facility includes two secondary redundancy appliances:
- a first secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in the storage associated with the secondary storage facility; and
 - a second secondary redundancy appliance shadowing the first secondary redundancy appliance, the second secondary redundancy appliance acting as the first secondary redundancy appliance in response to a failover event.
50. (New) The data redundancy method according to claim 48, wherein the primary storage facility includes two primary redundancy appliances:
- a first primary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in the storage associated with the primary storage facility; and
 - a second primary redundancy appliance shadowing the first primary redundancy appliance, the second primary redundancy appliance acting as the first primary redundancy appliance in response to a failover event.
51. (New) A data redundancy system, comprising:
- a primary storage facility including at least one primary redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility, the primary storage facility batching the write requests as send batches for transfer to a secondary storage facility; and
 - the secondary storage facility disjoint from the primary storage facility including at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility, the second storage facility batching the received write request as receive batches, the second redundancy appliance acting as the primary redundancy appliance in response to a failover event.
52. (New) The data redundancy system of claim 51, wherein the secondary storage facility includes two secondary redundancy appliances:
- a first secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in the storage associated with the secondary storage facility; and
 - a second secondary redundancy appliance shadowing the first secondary redundancy appliance, the second secondary redundancy appliance acting as the first secondary redundancy appliance in response to a failover event.

53. (New) The data redundancy system of claim 51, wherein the primary storage facility includes two primary redundancy appliances:
- a first primary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in the storage associated with the primary storage facility; and
 - a second primary redundancy appliance shadowing the first primary redundancy appliance, the second primary redundancy appliance acting as the first primary redundancy appliance in response to a failover event.
54. (New) The data redundancy system according to claim 51, wherein the failover event is selected from the group consisting of a manually initiated event and an automatically initiated event.
55. (New) The data redundancy system according to claim 51, wherein the failover event is a communication traffic condition.
56. (New) The data redundancy system according to claim 55, wherein the communication traffic condition includes a greater portion of request traffic originating closer to the second data storage facility.
57. (New) The data redundancy system according to claim 51, wherein a first write transaction in a send batch at the first storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction.
58. (New) The data redundancy system according to claim 51, wherein a first write transaction in a receive batch at the second storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction.
59. (New) The data redundancy system according to claim 51, wherein the primary redundancy appliance updates the secondary redundancy appliance with information about the data in storage associated with the primary storage facility by forwarding a write record to the secondary redundancy appliance for each write request, the secondary redundancy appliance sending acknowledgements to the primary redundancy appliance in batches in response to receiving write records.
60. (New) The data redundancy system according to claim 51, wherein the secondary redundancy appliance accesses data stored in the storage associated with the primary storage facility by the primary redundancy appliance when the secondary redundancy appliance acts as the primary redundancy appliance in response to the failover event.
61. (New) The data redundancy system according to claim 51, wherein a send batch size is adaptively determined based on available communication bandwidth between the primary storage facility and the secondary storage facility.

Remarks

The specification is amended to correct typographical errors. No new matter is introduced with the amendments to the specification presented herein.

Claims 1-47 are pending in the application. Claims 1, 3, 7, 10, 14, 16-19, 20, 21, 23, 25-30, 33, 34, 36, 38, 40, 41, 43-45 and 47 remain after amendment. Claims, 2, 4, 5, 6, 8, 9, 11, 12, 13, 15, 22, 24, 31, 32, 35, 37, 39, 42 and 46 are canceled. Claims 48-61 are new. No new matter is introduced with the amendments to the claims or with the addition of the new claims.

It is believed that the remarks laid out herein below attend to all rejections and further issues raised in the pending office action dated March 10, 2006.

Specification

The specification stands rejected based upon the current title. A more descriptive title is hereby proposed. Please replace the current title with the following:

DATA REDUNDANCY SYSTEM AND METHOD

Applicant has amended the specification to place it in accordance with the figures as originally filed with the application. No new matter is introduced with the amendments to the specification presented herein.

Applicant has corrected the reference number associated with "pending state" on page 23, line 12 "412" to read "**424**" such that the reference number is in accordance with Figure 7.

Applicant has corrected the reference number associated with "update state" on page 23, line 13 "414" to read "**430**" such that the reference number is in accordance with Figure 7.

Applicant has corrected the reference on page 27, line 27 "Figure 4" to read "**Figure 7**" such that the reference is in accordance with the figures as originally filed with the application.

Applicant has corrected the clause on page 30, lines 20-21 "[t]hough redundant logs may be maintained at each redundancy appliance the primary 102" to read "[t]hough redundant logs may be maintained at each redundancy appliance the in primary storage 102." to correct a grammatical error and to clarify the clause.

Applicant has corrected the reference number associated with "secondary facility" on page 33, line 18 "102" to read "**104**" such that the reference number is in accordance with Figure 1.

Applicant has corrected the reference number associated with "secondary facility" on page 35, line 2 "102" to read "104" such that the reference number is in accordance with Figure 1.

Applicant has corrected the reference number associated with "secondary", page 35, line 20 "102" to read "104" such that the reference number is in accordance with Figure 1.

Claims

Allowable subject matter - Claims 1, 3, 8, 9, 10, 11, 15 and 24-28

The Examiner has indicated that claims 9, 15, and 24-28 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims. Applicant respectfully thanks the Examiner for this identification of allowable subject matter, and accordingly makes the following amendments to the claims:

Applicant has amended claim 1 to add limitations of claim 24. Applicant has canceled claim 24 because it is incorporated into claim 1.

Applicant has amended claim 3 to add limitations of claims 8 and 9. Applicant has canceled claims 8 and 9 because they are incorporated into claim 3.

Applicant has amended claim 10 to incorporate the limitations of claims 11, 13, and 15. Applicant has canceled claims 11, 13, and 15 because they are incorporated into claim 10.

Applicant has amended claim 25 to properly depend from claim 1.

Applicant has amended claim 28 to change "second storage facility" to "secondary storage facility" in accordance with claim 1.

As a result of applicant's amendments, claims 1, 3, 10, and 25-28 each include matter the Examiner has identified as allowable. As such, Applicant respectfully requests withdrawal of the Examiner's rejections or objections.

Claim Rejections – 35 U.S.C. §102(b)

Claims 1-8, 10-12, 16-19, 29,32, 35-36 and 42-44 stand rejected under 35. U.S.C. §102(b) as being anticipated by U.S. Patent No: 6,694,447 to Leach et al., hereinafter "Leach". Applicants respectfully disagree.

In summary, the immediate application teaches a primary facility for storing a primary copy of data and a secondary facility for storing secondary copy of the data. Each storage facility has one or more redundancy appliances for storing data. The primary site updates the

secondary site using batches of data, wherein the data is buffered and transmitted to the secondary storage facility in batches. Data is buffered at the primary site until the batch terminates. Once terminated, the batch may be sent to the secondary site. The size of each batch may be determined upon available bandwidth; the more bandwidth, the smaller the batch size. Of note, the primary facility may include a first redundancy appliance and a second redundancy appliance; that is, the first and second redundancy appliances are within the primary facility. The secondary facility may include third and fourth redundancy appliances to match the redundancy appliances of the primary facility.

The primary and secondary facilities are capable of swapping rolls in the event of a failover and may remain in the swapped rolls or swap back upon full system recovery.

Leach, on the other hand, discloses a method and apparatus for increasing availability of an application during fail back from a secondary site to a primary site following a failure at the primary site. Leach defines fail back as a transition from the secondary site to the primary site after the primary site has been restored. Leach discloses a primary server for running an application located at the primary site and a secondary server located at a secondary site.

The method includes copying data from active storage volumes to secondary storage volumes on the secondary site while the application runs on the secondary site and updates the active storage volumes. Once the secondary storage volumes of the secondary site are updated, the data is re-synchronized from the secondary storage volumes on the secondary site to the primary storage volumes of the primary site. Following which, control is transferred back to the primary site, under the Leach fail back process. Leach does not disclose buffering or batching data between the primary site and the secondary site.

Independent Claim 1

With respect to amended claim 1, Leach does not disclose:

"a primary storage facility for storing a primary copy of data, **the primary storage facility including;**
 a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and
 a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance;
a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility **discrete from the primary storage facility and including;**
 a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and
 a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance."

Rather Leach teaches any and all redundancy systems as being present in a secondary site but not the primary site.

As Examiner indicated dependent claim 24 would be allowable if rewritten in independent form, the elements of claim 24 have been incorporated in independent claim 1. Moreover, claim 1 is now in allowable form as indicated by Examiner – clearly distinct from Leach.

Dependent claims 3, 7, 10, 14, 16-21 and 25-28

Claims 3, 7, 10, 14, 16-21 and 25-28 depend directly or indirectly from independent claim 1 in the re-written and allowable form, and thus are allowable for at least the same reasons, incorporated herein by reference. In addition, these dependent claims include additional limitations which are not included in Leach, such that Leach can no longer be said to anticipate them.

Claim 3 has been rewritten to include the indicated allowable subject matter of claim 9 through claim 8. Claim 10 has been rewritten to include the allowable subject matter of claim 15 claim 13, claim 12, and claim 11.

Claim 25-28 were object to by Examiner as being dependent upon a rejected claim. As claim 1 has been amended so as to now be allowable, claim 25 has been amended to properly depend from claim 1. Claims 26-28 depend from claim 25 and are thus now also in proper form for allowance.

As such applicant respectfully requests withdrawal of Examiners rejections, and allowance of claims 1, 3, 7, 10, 14, 16-21 and 25-28.

Independent Claim 29

With respect to claim 29, Leach does not disclose:

receiving a sequence of write requests at a first redundancy appliance at the primary storage facility;
storing data for the sequence of write requests in primary storage at the primary storage facility;
logging at least a sequence of write requests for transfer to the secondary storage facility;
forwarding the logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility ***batching the received write request as receive batches;***
storing data associated with the receive batches in secondary storage at the secondary storage facility; and
the second redundancy appliance acting as the first redundancy appliance in response to a failover event.

As noted above, Leach does not disclose sending batches between the primary and secondary facility, and therefore cannot anticipate claim 29.

Dependent claims 30, 33, 34, 36, 38, 40, 41, 43, 44 and 45

Claims 30, 33, 34, 36, 38, 40, 41, 43, 44 and 45 depend directly or indirectly from independent claim 29, and thus benefit from like argument, incorporated herein by reference.

As such applicant respectfully requests withdrawal of Examiners rejections, and allowance of claims 29, 30, 33, 34, 36, 38, 40, 41, 43, 44 and 45.

Claim Rejections – 35 U.S.C. §103(a)

To establish a *prima facie* case of obviousness under 35 U.S.C. § 103(a), three basic criteria must be met:

(a) **there must be some suggestion or motivation**, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the references or to combine reference teachings;

(b) there must be a **reasonable expectation of success**; *and*

(c) **the prior art reference** (or references when combined) ***must teach or suggest all the claim limitations.*** The ***teaching or suggestion*** to make the claimed combination ***and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure...*** (emphasis added) MPEP § 2143.

"The mere fact that the prior art *may* be modified in the manner suggested by the Examiner *does not make the modification obvious* unless the prior art suggest the desirability of the modification." (emphasis added) *In re Fritch*, 23 USPQ2d 1780,1783-84 (Fed. Cir. 1992).

Claims 13-14 and 33-34

Claims 13-14 and 33-34 stand rejected under 35 U.S.C. §103(a) as being unpatentability anticipated by U.S. Patent No: 5,889,935 to Ofek et al., hereinafter "Ofek". Applicant respectfully disagrees.

Ofek discloses a data processing system which controls storing of data received from a host on a primary data storage system and copying of the data to a remote secondary data storage system. The primary data storage system is connected to the remote secondary data storage system by at least one link. Each storage system includes a system controller and a

storage device. One or more of the system controllers coordinates copying of the data from the primary data storage system to the secondary data storage system. At least one of the system controllers maintains a list the data that is to be copied to the secondary data storage system. Data may be transferred between the two storage systems in a synchronous mode, a semi-synchronous mode, an adaptive copy – write pending mode, and an adaptive copy – disk mode. Ofek does not disclose buffering or batching data between the primary site and the secondary site.

Claim 13 has been canceled without prejudice and is therefore no longer at issue. Claim 14 has been amended to depend from claim 1 in rewritten allowable form through dependent claim 10, and is therefore allowable for at least the same reasons stated above with respect to claims 1 and 10, and incorporated herein by reference.

Claims 33 and 34 depend from claim 29 which in amended form discloses forwarding the logged write requests in batches. Neither Leach nor Ofek teach or suggest communication between the primary and secondary locations occurring in batches. As such, whether separately or combined, Leah and Ofek fail to render either claim 33 or 34 as obvious. In addition, there is neither motivation to combine the references, nor an implied or express expectation of success in making such a combination, or any indication that such a combination would result in the batch style communication as set forth by applicant. Withdrawal of Examiner's rejection and allowance of claims 14 and 33-34 is therefore respectfully requested.

Claims 20-23, 37-41 and 45-57

Claims 20-23, 37-41 and 45-57 stand rejected under 35 U.S.C. §103(a) as being unpatentability anticipated by U.S. Patent No: 6,442,706 to Wahl et al., hereinafter "Wahl". Applicant respectfully disagrees.

Wahl discloses a system for remote mirroring of data including a primary system and a secondary system. The primary system runs applications and provides data storage services during normal operation. The secondary system stores a copy of data stored on the primary system. The primary system includes at least one local storage device to store data and at least one write log device for journaling data written to the at least one local storage device. The secondary system includes at least one mirror device to mirror data stored on the one or more local storage device. The primary system and secondary system each include a mirror daemon which control the transfer of data between systems. During normal operation, the primary mirror daemon sends data one at a time across the network to the remote mirror daemon (Col 19, lines 56-60). Wahl does not disclose sending batches between the primary and secondary facility. Further, Wahl does not teach redundant appliances mirroring each other within the primary data facility and redundant appliances mirroring each other within the secondary data facility.

Claim 22 has been canceled without prejudice and is therefore no longer at issue. Claims 20, 21 and 23 depended directly and indirectly from claim 1, now appearing in amended rewritten form for allowance. As such claims 20, 21 and 23 are allowable for at least the same reasons as claim 1, incorporated herein by reference.

Claims 37, 39 and 46 have been canceled without prejudice and are therefore no longer at issue. Claims 38, 40, 41, 45, and 47 depend directly or indirectly from independent claim 29, which discloses sending send batches which in turn are received and stored as receive batches, operations neither taught or suggested by Wahl. As such, claims 38, 40, 41, 45 and 47 are allowable for at least the same reasons as claim 29, incorporated herein by reference.

With respect to claim 40, Wahl discusses adjusting bandwidth usage based on user definable parameters, but does not teach or suggest adjusting the bandwidth usage directly based upon the bandwidth availability. As Wahl does not disclose communication with batches, it cannot teach or suggest adaptively adjusting batch size based on available bandwidth.

With respect to claim 41, respectfully Wahl does not disclose replacing a first write transaction in a send batch at the primary storage facility with a second write transaction where the second write transaction overwrites data written by the first write transaction batch. In Col 7, lines 48-53 Wahl teaches replacing older data entries with new data entries **only after** there is confirmation that the older data entry has been committed. By such a method all entries must be committed, whereas applicant teaches recognizing and consolidating actions within a batch. This advantageously reduces the instructions sets within the batch, and eliminates redundant processor execution time to execute repeated instructions upon the same data element.

With respect to claim 45, respectfully Wahl does not disclose determining whether a receive batch is safe to apply before storing the receive batch at the secondary storage facility by performing a data integrity test on the data of the receive batch.

With respect to claim 47, respectfully Wahl does not disclose batching write requests received at secondary storage facility as receive batches. Furthermore, Wahl does not disclose replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction.

As such, whether separately or combined, Leach and Wahl fail to render either claim 38, 40, 41, 45, or 47 as obvious. In addition, there is neither motivation to combine the references, nor an implied or express expectation of success in making such a combination, or any indication that such a combination would result in the batch style communication as set forth by applicant, communication based on the available bandwidth itself, or the identification

of and consolidation of repetitious transactions prior to execution at the secondary data facility.

Withdrawal of Examiner's rejection and allowance of claims 20, 21 23, 38, 40, 41, 45 and 47 is therefore respectfully requested.

New Claims

In light of Examiner's comments and the cited art, Applicant presents new claims 48-61. The remarks presented herein above apply equally to new claims 48-61. For example, claims 48-50 depend from claim 29 and benefit from like argument, incorporated herein by reference.

Claims 51-61 distinctly draw out and identify the nature of the batched communications (i.e., send batches and receive batches), that send batch size can be adaptively determined by available bandwidth, and that redundant write actions can be consolidated for efficiency, as well as other patentably distinct elements neither taught, described or implied by Leach, Ofek or Wahl whether viewed separately or collectively.

Canceled Claims

Applicant has cancel claims 2, 4, 5, 6, 8, 9, 11, 12, 13, 15, 22, 24, 31, 32, 35, 37, 39, 42, 46 without prejudice.

Conclusion

In view of the above Remarks and the amendments to the specification and claims, Applicant has addressed all issues raised in the Office Action dated March 10, 2006, and respectfully solicits a Notice of Allowance for claims 1, 3, 7, 10, 14, 16-19, 20, 21, 23, 25-30, 33, 34, 36, 38, 40, 41, 43, 44, 45, and 47-61. Should any issues remain, the Examiner is encouraged to telephone the undersigned attorney.

It is believed that all of the pending claims have been addressed. However, the absence of a reply to a specific rejection, issue or comment does not signify agreement with or concession of that rejection, issue or comment. In addition, because the arguments made above may not be exhaustive, there may be reasons for patentability of any or all pending claims (or other claims) that have not been expressed. Finally, nothing in this paper should be construed as an intent to concede any issue with regard to any claim, except as

specifically stated in this paper, and the amendment of any claim does not necessarily signify concession of unpatentability of the claim prior to its amendment.

Applicant originally paid for 47 claims, including 3 independent claims. As a result of applicant's amendments, the application now includes 42 claims, 3 of which are independent claims. Consequently, applicant believes that no fees are due; however, should any fee be deemed necessary in connection with this Amendment and Response, the Commissioner is authorized to charge deposit account 08-2025, referencing the Attorney Docket Number 200309044-1.

Respectfully submitted,

By:



Daniel W. Roberts,
Reg. No. 52,172
LATHROP & GAGE L.C.
4845 Pearl East Circle, Suite 300
Boulder, CO 80301
Telephone: (720) 931-3016
Facsimile: (720) 931-3001

PATENT APPLICATION FEE DETERMINATION RECORD
Effective December 8, 2004

10/456053

CLAIMS AS FILED - PART I

	(Column 1)	(Column 2)
TOTAL CLAIMS		
FOR	NUMBER FILED	NUMBER EXTRA
TOTAL CHARGEABLE CLAIMS	47 minus 20 =	27
INDEPENDENT CLAIMS	2 minus 3 =	0
MULTIPLE DEPENDENT CLAIM PRESENT		☐

* If the difference in column 1 is less than zero, enter "0" in column 2

CLAIMS AS AMENDED - PART II

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT A	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
	6-504		
Total	* 42	Minus ** 47	= 0
Independent	* 3	Minus *** 3	= 0
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM			☐

SMALL ENTITY TYPE ☐ OR

OTHER THAN SMALL ENTITY

RATE	FEE
BASIC FEE	
X\$ 25=	
X100=	
+180=	
TOTAL	

RATE	FEE
BASIC FEE	750
X\$ 25=	486
X100=	
+180=	
TOTAL	1236

SMALL ENTITY OR

OTHER THAN SMALL ENTITY

RATE	ADDITIONAL FEE
X\$ 25=	
X100=	
+180=	
TOTAL ADDIT. FEE	

RATE	ADDITIONAL FEE
X\$ 25=	
X100=	
+180=	
TOTAL ADDIT. FEE	

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT B	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
Total	*	Minus **	=
Independent	*	Minus ***	=
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM			☐

RATE	ADDITIONAL FEE
X\$ 25=	
X100=	
+180=	
TOTAL ADDIT. FEE	

RATE	ADDITIONAL FEE
X\$ 25=	
X100=	
+180=	
TOTAL ADDIT. FEE	

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT C	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
Total	*	Minus **	=
Independent	*	Minus ***	=
FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM			☐

RATE	ADDITIONAL FEE
X\$ 25=	
X100=	
+180=	

RATE	ADDITIONAL FEE
X\$ 25=	
X100=	
+180=	



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/456,053	06/06/2003	Minwen Ji	200309044-1	5067

22879 7590 03/10/2006

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER

MCCARTHY, CHRISTOPHER S

ART UNIT PAPER NUMBER

2113

DATE MAILED: 03/10/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No. 10/456,053	Applicant(s) JI ET AL.	
	Examiner Christopher S. McCarthy	Art Unit 2113	
	-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --		

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) ☒ Responsive to communication(s) filed on 06 June 2003.

2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.

3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) ☒ Claim(s) 1-47 is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5) ☐ Claim(s) _____ is/are allowed.

6) ☒ Claim(s) 1-8, 10-14, 16-23 and 29-47 is/are rejected.

7) ☒ Claim(s) 9, 15 and 24-28 is/are objected to.

8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) ☒ The specification is objected to by the Examiner.

10) ☒ The drawing(s) filed on 6/6/03 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 a) ☐ All b) ☐ Some * c) ☐ None of:
 1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892) 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948) 3) ☒ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date <u>11/21/03</u> .	4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____. 5) ☐ Notice of Informal Patent Application (PTO-152) 6) ☐ Other: _____.
--	--

DETAILED ACTION

Specification

The title of the invention is not descriptive. A new title is required that is clearly indicative of the invention to which the claims are directed.

Claim Objections

Claim 8 is objected to because of the following informalities: The word “condition” is misspelled. Appropriate correction is required.

Double Patenting

Claim 10 is objected to under 37 CFR 1.75 as being a substantial duplicate of claim 2. When two claims in an application are duplicates or else are so close in content that they both cover the same thing, despite a slight difference in wording, it is proper after allowing one claim to object to the other as being a substantial duplicate of the allowed claim. See MPEP § 706.03(k).

Claim Rejections - 35 USC § 112

The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claim 30 recites the limitation "said storing and arranging" in the claim language. There is insufficient antecedent basis for this limitation in the claim. "Arranging" is not a prior disclosed step.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1-8, 10-12, 16-19, 29-32, 35-36, 42-44 are rejected under 35 U.S.C. 102(e) as being anticipated by Leach et al. U.S. Patent 6,694,447.

As per claim 1, Leach teaches a data redundancy system including a primary facility for storing a primary copy of data and a secondary facility for storing data that is redundant of the primary copy of the data, wherein the primary facility comprises: a first redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance (column 6, lines 19-35).

As per claim 2, Leach teaches the data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage (column 6, lines 19-35).

As per claim 3, Leach teaches the data redundancy system according to claim 1, wherein the second redundancy appliance assumes the role of the first redundancy appliance in response to a failover event (column 5, lines 16-18).

As per claim 4, Leach teaches the data redundancy system according to claim 3, wherein the failover event comprises a fault (column 5, lines 16-18).

As per claim 5, Leach teaches the data redundancy system according to claim 4, wherein the fault occurs at a location selected from the group consisting of the first storage facility, the second storage facility and a communication medium between the first data storage facility and the second data storage facility (column 5, lines 16-18).

As per claim 6, Leach teaches the data redundancy system according to claim 3, wherein the failover event comprises an operational event (column 6, line 56 – column 7, line 2).

As per claim 7, Leach teaches the data redundancy system according to claim 6, wherein the failover event is selected from the group consisting of a manually initiated event and an automatically initiated event (column 6, lines 56-63).

As per claim 8, Leach teaches the data redundancy system according to claim 3, wherein the failover event comprises a communication traffic condition (column 6, lines 57-60, wherein a server failure would be a communication condition).

As per claim 10, Leach teaches the data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage (column 6, lines 44-56).

As per claim 11, Leach teaches the data redundancy system according to claim 10, wherein the first redundancy appliance updates the second redundancy appliance by forwarding a write record to the second redundancy appliance for each write request (column 6, lines 19-35).

As per claim 12, Leach teaches the data redundancy system according to claim 11, wherein write records are forwarded to the second appliance in batches (column 6, lines 19-35).

As per claim 16, Leach teaches the data redundancy system according to claim 1, wherein the first redundancy appliance is disjoint from the second redundancy appliance (column 4, lines 16-37).

As per claim 17, Leach teaches the data redundancy system according to claim 1, wherein the second redundancy appliance maintains a shadow copy of data stored by the first redundancy appliance wherein the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance (column 7, lines 23-39).

As per claim 18, Leach teaches the data redundancy system according to claim 1, wherein the second redundancy appliance accesses the data stored in storage by the first redundancy appliance when the second redundancy appliance assumes the role of the first redundancy appliance (column 7, lines 2-5).

As per claim 19, Leach teaches the data redundancy system according to claim 1, wherein the first redundancy appliance forwards redundant data to the secondary facility according to send batches (column 6, lines 19-35).

As per claim 29, Leach teaches a data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility, the method comprising: receiving a sequence of write requests at a first redundancy appliance at the primary storage facility; storing data for the sequence of write requests in storage at the primary storage facility; updating a shadow redundancy appliance at the primary storage facility with information about the data in storage; and forwarding the sequence of write requests to the secondary storage facility (column 6, lines 19-35; column 5, lines 12-15).

As per claim 30, Leach teaches the data redundancy method according to claim 29, wherein said first redundancy appliance performs said storing and arranging prior to a failover event and wherein the second redundancy appliance performs said storing and arranging after the failover event (column 6, lines 19-35).

As per claim 31, Leach teaches the data redundancy method according to claim 29, wherein said updating is performed synchronously (column 6, lines 19-35).

As per claim 32, Leach teaches the data redundancy method according to claim 31, wherein said updating comprises the first redundancy appliance forwarding a write record to the second redundancy appliance for each write request (column 6, lines 19-35).

As per claim 35, Leach teaches the data redundancy method according to claim 29, wherein the first redundancy appliance is disjoint from the second redundancy appliance (column 4, lines 16-49).

As per claim 36, Leach teaches the data redundancy method according to claim 29, further comprising the second redundancy appliance maintaining a shadow copy of data stored by the first redundancy appliance (column 7, lines 23-29).

As per claim 42, Leach teaches the data redundancy method according to claim 29, further comprising: receiving write requests from the primary storage facility at a third redundancy appliance at the secondary storage facility; storing data for the write requests in storage at the secondary storage facility (column 6, lines 19-35).

As per claim 43, Leach teaches the data redundancy method according to claim 42, further comprising updating a shadow redundancy appliance at the secondary storage facility with information about the data in storage at the secondary storage facility (column 6, lines 19-35).

As per claim 44, Leach teaches the data redundancy method according to claim 29, further comprising storing data for the write requests in storage at the secondary storage facility according to receive batches (column 6, lines 19-35).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 13-14, 33-34 are rejected under 35 U.S.C. 103(a) as being unpatentable over Leach in view of Ofek et al. U.S. Patent 5,889,935.

As per claim 13, Leach teaches the data redundancy system according to claim 11. Leach does not explicitly teach wherein the second redundancy appliance sends an acknowledgement in response to receiving the write record. Ofek does teach wherein the second redundancy appliance sends an acknowledgement in response to receiving the write record (column 3, lines 28-37). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the acknowledgement system of Ofek in the failover system of Leach. One of ordinary skill in the art would have been motivated to use the acknowledgement system of Ofek in the failover system of Leach because Ofek teaches the synchronizing of primary data to replicated, backup storage; an explicit desire of Leach (abstract).

As per claim 14, Leach teaches the data redundancy system according to claim 13. Leach does not explicitly teach wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance. Ofek does teach wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance (column 3, lines 28-37). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the acknowledgement system of Ofek in the failover system of Leach. One of ordinary skill in the art would have been motivated to use the acknowledgement system of Ofek in the failover system of Leach because Ofek teaches the synchronizing of primary data to replicated, backup storage; an explicit desire of Leach (abstract).

As per claim 33, Leach teaches the data redundancy method according to claim 32. Leach does now teach the the second redundancy appliance sending an acknowledgement in response to receiving the write record. Ofek does teach the second redundancy appliance sending an acknowledgement in response to receiving the write record (column 3, lines 28-37). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the acknowledgement system of Ofek in the failover system of Leach. One of ordinary skill in the art would have been motivated to use the acknowledgement system of Ofek in the failover system of Leach because Ofek teaches the synchronizing of primary data to replicated, backup storage; an explicit desire of Leach (abstract).

As per claim 34, Leach teaches the data redundancy method according to claim 33. Leach does not teach the first redundancy appliance sending an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance. Ofek does teach the first redundancy appliance sending an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance (column 3, lines 28-37). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the acknowledgement system of Ofek in the failover system of Leach. One of ordinary skill in the art would have been motivated to use the acknowledgement system of Ofek in the failover system of Leach because Ofek teaches the synchronizing of primary data to replicated, backup storage; an explicit desire of Leach (abstract).

Claims 20-23, 37-41, 45-47 are rejected under 35 U.S.C. 103(a) as being unpatentable over Leach in view of Wahl et al. U.S. Patent 6,442,706.

As per claim 20, Leach teaches the data redundancy system according to claim 1. Leach does not teach wherein a first write transaction in a send batch at the first storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction. Wahl does teach wherein a first write transaction in a send batch at the first storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction (column 7, lines 48-53). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 21, Leach teaches the data redundancy system according to claim 20. Leach does not teach wherein a send batch size is determined according to a time interval for collecting transactions into the send batch. Wahl does teach wherein a send batch size is determined according to a time interval for collecting transactions into the send batch (column 17, lines 11-27). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the

Art Unit: 2113

primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 22, Leach teaches the data redundancy system according to claim 20. Leach does not teach wherein a send batch size is determined according to a predetermined quantity of transactions to be included in the send batch. Wahl does teach a send batch size is determined according to a predetermined quantity of transactions to be included in the send batch (column 17, line 64 – column 18, line 14). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 23, Leach teaches the data redundancy system according to claim 20. Leach does not teach wherein a send batch size is determined adaptively based on available communication bandwidth between the first storage facility and the second storage facility. Wahl teaches a send batch size is determined adaptively based on available communication bandwidth between the first storage facility and the second storage facility (column 14, lines 53-55). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the

Art Unit: 2113

primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 37, Leach teaches the data redundancy method according to claim 29.

Leach does not teach further comprising the first redundancy appliance arranging the sequence of write requests according to a sequence of send batches, wherein said forwarding comprises forwarding the send batches to the secondary facility. Wahl does the first redundancy appliance arranging the sequence of write requests according to a sequence of send batches, wherein said forwarding comprises forwarding the send batches to the secondary facility (column 7, lines 48-51; column 17, lines 11-27, line 64 – column 18, line 14). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 38, Leach teaches the data redundancy method according to claim 37.

Leach does not teach further comprising determining a send batch size according to a time interval for collecting transactions into the send batch. Wahl teaches determining a send batch size according to a time interval for collecting transactions into the send batch (column 17, lines 11-27). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover

system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 39, Leach teaches the data redundancy method according to claim 37. Leach does not teach further comprising determining a send batch size according to a predetermined quantity of transactions to be included in the send batch. Wahl teaches determining a send batch size according to a predetermined quantity of transactions to be included in the send batch (column 17, line 64 – column 18, line 14). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 40, Leach teaches the data redundancy method according to claim 37. Leach does not teach further comprising determining a send batch size adaptively based on available communication bandwidth between the first storage facility and the second storage facility. Wahl does teach determining a send batch size adaptively based on available communication bandwidth between the first storage facility and the second storage facility (column 14, lines 53-55).

As per claim 41, Leach teaches the data redundancy method according to claim 29. Leach does not teach further comprising replacing a first write transaction in a send batch at the

Art Unit: 2113

first storage facility with a second write transaction where the second write transaction overwrites data written by the first write transaction. Wahl teaches replacing a first write transaction in a send batch at the first storage facility with a second write transaction where the second write transaction overwrites data written by the first write transaction (column 7, lines 48-53). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 45, Leach teaches the data redundancy method according to claim 44. Leach does not teach further comprising determining whether a receive batch is safe to apply before storing the receive batch. Wahl teaches determining whether a receive batch is safe to apply before storing the receive batch (column 24, lines 34-42). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 46, Leach teaches the data redundancy method according to claim 45. Leach does not teach wherein said determining whether a receive batch is safe to apply

Art Unit: 2113

comprises performing a data integrity test on data of the receive batch. Wahl teaches determining whether a receive batch is safe to apply comprises performing a data integrity test on data of the receive batch (column 24, lines 24-42). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

As per claim 47, Leach teaches the data redundancy method according to claim 44. Leach does not teach further comprising replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction. Wahl teaches replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction (column 7, lines 48-53). It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the failover system of Wahl to the failover system of Leach. One of ordinary skill in the art would have been motivated to use the failover system of Wahl to the failover system of Leach because Wahl teaches the synchronized replication of a secondary device to backup the primary device upon failure (column 3, lines 14-19; column 4, lines 58-62); an explicit desire of Leach (column 6, lines 22-32, 57-60).

Allowable Subject Matter

Claims 9, 15, 24-28 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

Conclusion

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure: See attached PTO-892.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Christopher S. McCarthy whose telephone number is (571)272-3651. The examiner can normally be reached on M-F, 9 - 5:30.


If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Robert Beausoliel can be reached on (571)272-3645. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2113

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

csm
March 3, 2006


ROBERT BEAUSOLIEL
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100



PATENT APPLICATION

Sheet 1 of 3

FORM PTO-1449

LIST OF PATENTS AND PUBLICATIONS FOR
APPLICANT'S INFORMATION DISCLOSURE
STATEMENT

(Use several sheets if necessary)

ATTY. DOCKET NO.

200309044-1

APPLICATION NO.

10/456,053

CONFIRMATION NO.

APPLICANT

Ji et al.

FILING DATE

06/06/03

GROUP

REFERENCE DESIGNATION

U.S. PATENT DOCUMENTS

EXAMINER INITIAL		DOCUMENT NUMBER	PUBLICATION DATE	NAME	Pages, Columns, Lines Where Relevant Passages or Figures Appear
<i>me</i>	1A	5,544,347	Aug. 6, 1996	Yanai et al.	
	1B	5,742,792	Apr. 21, 1998	Yanai et al.	
	1C	5,909,692	Jun. 1, 1999	Yanai et al.	
	1D	6,092,066	Jul. 18, 2000	Ofek	
	1E	6,101,497	Aug. 8, 2000	Ofek	
	1F	6,108,748	Aug. 22, 2000	Ofek et al.	
	1G				
	1H				
	1I				
	1J				
	1K				

FOREIGN PATENT DOCUMENTS

		DOCUMENT NUMBER	PUBLICATION DATE	NAME OF PATENTEE OR APPLICANT	Pages/Columns/Lines Where Relevant Passages/Figures Appear	Check if Translation attached
	1L					
	1M					
	1N					
	1O					
	1P					

OTHER REFERENCES (Including Author, Title, Date, Pertinent Pages, etc.)

<i>me</i>	1Q	Fay Chang, Minwen Ji, Shun-Tak A. Leung, John MacCormick, Sharon E. Perl and Li Zhang, Myriad: Cost-effective Disaster Tolerance, Proceedings of the FAST 2002 Conference on File and Storage Technologies, Monterey, CA, pp. 103-116, The USENIX Association, Berkeley, CA, January 2002.
<i>me</i>	1R	Chia Chao, Robert English, David Jacobson, Alexander Stepanov, and John Wilkes, Mime: a high performance parallel storage device with strong recovery guarantees, HP Laboratories Technical Report HPL-CSP-92-9 rev 1, 18 March 1992, revised 6 November 1992, Hewlett-Packard Company, Palo Alto, CA, 1992.
	1S	~~Remote Copy Administrator's Guide and Reference, IBM DFSMS/MVS Version 4, Fourth Edition, Pages i to xv and pages 1 thru 170, International Business Machines Corporation, Raleigh, NC, December 1997.~~

EXAMINER

DATE CONSIDERED

2-9-06

Rev 02/03 (PTO1449)



PATENT APPLICATION

Sheet 2 of 3

FORM PTO-1449

LIST OF PATENTS AND PUBLICATIONS FOR
APPLICANT'S INFORMATION DISCLOSURE
STATEMENT

(Use several sheets if necessary)

ATTY. DOCKET NO.

200309044-1

APPLICATION NO.

10/456,053

CONFIRMATION NO.

APPLICANT

Ji et al.

FILING DATE

06/06/03

GROUP

REFERENCE DESIGNATION

U.S. PATENT DOCUMENTS

EXAMINER INITIAL	DOCUMENT NUMBER	PUBLICATION DATE	NAME	Pages, Columns, Lines Where Relevant Passages or Figures Appear
	2A			
	2B			
	2C			
	2D			
	2E			
	2F			
	2G			
	2H			
	2I			
	2J			
	2K			

FOREIGN PATENT DOCUMENTS

	DOCUMENT NUMBER	PUBLICATION DATE	NAME OF PATENTEE OR APPLICANT	Pages/Columns/Lines Where Relevant Passages/Figures Appear	Check if Translation attached
	2L				
	2M				
	2N				
	2O				
	2P				


OTHER REFERENCES (Including Author, Title, Date, Pertinent Pages, etc.)

2Q	Hugo Patterson, Stephen Manley, Mike Federwisch, Dave Hitz, Steve Kleiman and Shane Owara, SnapMirror: File System Based Asynchronous Mirroring for Disaster Recovery, Proceedings of the FAST 2002 Conference on File and Storage Technologies, Monterey, CA, pp. 117-129, The USENIX Association, Berkeley, CA, January 2002.
2R	Symmetrix Remote Data Facility (SRDF) Product Description Guide, EMC Corporation, Hopkinton, MA, USA, 2000.
2S	Michael Stonebraker and Gerhard A. Schloss, Distributed Raid - A New Multiple Copy Algorithm, UCB Electronics Research Laboratory Technical Report ERL-M89-56, Electronics Research Laboratory, University of California, Berkeley, CA, 1989.

EXAMINER

DATE CONSIDERED

Rev 02/03 (PTO1449)



PATENT APPLICATION

Sheet 3 of 3

FORM PTO-1449

LIST OF PATENTS AND PUBLICATIONS FOR
APPLICANT'S INFORMATION DISCLOSURE
STATEMENT

(Use several sheets if necessary)

ATTY. DOCKET NO.

200309044-1

APPLICATION NO.

10/456,053

CONFIRMATION NO.

APPLICANT

Ji et al.

FILING DATE

06/06/03

GROUP

REFERENCE DESIGNATION

U.S. PATENT DOCUMENTS

EXAMINER INITIAL	DOCUMENT NUMBER	PUBLICATION DATE	NAME	Pages, Columns, Lines Where Relevant Passages or Figures Appear
	3A			
	3B			
	3C			
	3D			
	3E			
	3F			
	3G			
	3H			
	3I			
	3J			
	3K			

FOREIGN PATENT DOCUMENTS

	DOCUMENT NUMBER	PUBLICATION DATE	NAME OF PATENTEE OR APPLICANT	Pages/Columns/Lines Where Relevant Passages/Figures Appear	Check if Translation attached
	3L				
	3M				
	3N				
	3O				
	3P				

OTHER REFERENCES (including Author, Title, Date, Pertinent Pages, etc.)

3Q	EMC TimeFinder Product Description Guide, EMC Corporation, Hopkinton, Massachusetts, USA, 1998.
3R	
3S	

EXAMINER

DATE CONSIDERED

2-9-06

Rev 02/03 (PTO-1449)

Notice of References Cited	Application/Control No. 10/456,053	Applicant(s)/Patent Under Reexamination JI ET AL.	
	Examiner Christopher S. McCarthy	Art Unit 2113	Page 1 of 2

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-6,694,447 B1	02-2004	Leach et al.	714/6
*	B	US-6,148,383 A	11-2000	Micka et al.	711/162
*	C	US-5,889,935 A	03-1999	Ofek et al.	714/6
*	D	US-6,442,706 B1	08-2002	Wahl et al.	714/6
*	E	US-6,898,685 B2	05-2005	Meiri et al.	711/167
*	F	US-6,769,030 B1	07-2004	Bournas, Redha M.	709/233
*	G	US-2003/0074600 A1	04-2003	Tamatsu, Masaharu	714/6
*	H	US-6,157,991 A	12-2000	Arnon, Dan	711/161
*	I	US-6,487,561 B1	11-2002	Ofek et al.	707/204
*	J	US-6,615,332 B2	09-2003	Yamamoto et al.	711/167
*	K	US-2003/0014534 A1	01-2003	Watanabe et al.	709/233
*	L	US-6,442,709 B1	08-2002	Beal et al.	714/33
*	M	US-2004/0088507 A1	05-2004	Satoyama et al.	711/162

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Notice of References Cited	Application/Control No. 10/456,053	Applicant(s)/Patent Under Reexamination JI ET AL.	
	Examiner Christopher S. McCarthy	Art Unit 2113	Page 2 of 2

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-6,728,848 B2	04-2004	Tamura et al.	711/162
*	B	US-2003/0225760 A1	12-2003	Ruuth et al.	707/5
*	C	US-6,591,351 B1	07-2003	Urabe et al.	711/162
*	D	US-5,140,592 A	08-1992	Idleman et al.	714/5
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Search Notes

Application/Control No.

10/456,053

Examiner

Christopher S. McCarthy

Applicant(s)/Patent under
Reexamination

JI ET AL.

Art Unit

2113

SEARCHED

Class	Subclass	Date	Examiner
718	101	2/16/2006	CM
714	6	2/13/2006	CM
709	233	2/16/2006	CM
711	162	2/15/2006	CM

INTERFERENCE SEARCHED

Class	Subclass	Date	Examiner

**SEARCH NOTES
(INCLUDING SEARCH STRATEGY)**

	DATE	EXMR
714/43 - text search only - see attached search report	2/10/2006	CM
Search History Enclosed	3/3/2006	CM

Index of Claims 		Application/Control No. 10/456,053 Examiner Christopher S. McCarthy		Applicant(s)/Patent under Reexamination JI ET AL. Art Unit 2113	
---	--	--	--	--	--

√ Rejected = Allowed	- (Through numeral) Cancelled + Restricted	N Non-Elected I Interference	A Appeal O Objected
---------------------------------------	---	---	--------------------------------------

Claim		Date			
Final	Original	3/3/06			
1	√				
2	√				
3	√				
4	√				
5	√				
6	√				
7	√				
8	√				
9	o				
10	√				
11	√				
12	√				
13	√				
14	√				
15	o				
16	√				
17	√				
18	√				
19	√				
20	√				
21	√				
22	√				
23	√				
24	o				
25	o				
26	o				
27	o				
28	o				
29	√				
30	√				
31	√				
32	√				
33	√				
34	√				
35	√				
36	√				
37	√				
38	√				
39	√				
40	√				
41	√				
42	√				
43	√				
44	√				
45	√				
46	√				
47	√				
48					
49					
50					

Claim		Date			
Final	Original				
51					
52					
53					
54					
55					
56					
57					
58					
59					
60					
61					
62					
63					
64					
65					
66					
67					
68					
69					
70					
71					
72					
73					
74					
75					
76					
77					
78					
79					
80					
81					
82					
83					
84					
85					
86					
87					
88					
89					
90					
91					
92					
93					
94					
95					
96					
97					
98					
99					
100					

Claim		Date			
Final	Original				
101					
102					
103					
104					
105					
106					
107					
108					
109					
110					
111					
112					
113					
114					
115					
116					
117					
118					
119					
120					
121					
122					
123					
124					
125					
126					
127					
128					
129					
130					
131					
132					
133					
134					
135					
136					
137					
138					
139					
140					
141					
142					
143					
144					
145					
146					
147					
148					
149					
150					



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
 United States Patent and Trademark Office
 Address: COMMISSIONER FOR PATENTS
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
 www.uspto.gov



Bib Data Sheet

CONFIRMATION NO. 5067

SERIAL NUMBER 10/456,053	FILING DATE 06/06/2003 RULE	CLASS 714	GROUP ART UNIT 2113	ATTORNEY DOCKET NO. 200309044-1		
APPLICANTS Minwen Ji, Sunnyvale, CA; <i>MJ</i> Alistair Veitch, Mountain View, CA; John Wilkes, Palo Alto, CA;						
** CONTINUING DATA *****						
** FOREIGN APPLICATIONS *****						
IF REQUIRED, FOREIGN FILING LICENSE GRANTED ** 07/29/2003						
Foreign Priority claimed ☐ yes ☒ no 35 USC 119 (a-d) conditions ☐ yes ☒ no ☐ Met after met Verified and Acknowledged	Allowance <i>[Signature]</i> Examiner's Signature	Initials <i>[Initials]</i>	STATE OR COUNTRY CA	SHEETS DRAWING 14	TOTAL CLAIMS 47	INDEPENDENT CLAIMS 2
ADDRESS 22879 HEWLETT PACKARD COMPANY P O BOX 272400, 3404 E. HARMONY ROAD INTELLECTUAL PROPERTY ADMINISTRATION FORT COLLINS , CO 80527-2400						
TITLE Fault-tolerant data redundancy technique						
FILING FEE RECEIVED 1366	FEES: Authority has been given in Paper No. _____ to charge/credit DEPOSIT ACCOUNT No. _____ for following:			☐ All Fees ☐ 1.16 Fees (Filing) ☐ 1.17 Fees (Processing Ext. of time) ☐ 1.18 Fees (Issue)		

<http://neo:8000/PrexServlet/PrexAction?serviceName=BibDataSheet&Action=display&brows...> 3/3/06

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	6	((("5544347") or ("5742792") or ("5909692") or ("6092066") or ("6101497") or ("6108748"))).PN.	USPAT	OR	OFF	2006/02/09 15:51
S2	372	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:21
S14	372	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:16
S15	0	718/101.ccls. and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:17
S16	11	718/101.ccls. and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:17
S18	85	714/6.ccls. and batch	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:03
S19	94	714/6.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 11:22
S20	0	("2004/0268177").URPN.	USPAT	OR	OFF	2006/02/10 11:24
S21	721	714/43.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 12:33
S22	23	714/43.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 12:37
S23	87	714/4.ccls. and batch\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 12:56
S24	13260	((batch\$3 or transaction or packet) with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S25	400	((batch\$3 or transaction or packet) with adjust\$4 with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S26	370	((batch\$3 or packet) with adjust\$4 with size) and (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 12:57
S27	26	((batch\$3 or packet) with adjust\$4 with size) same (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 13:11
S28	0	("2004/0260977").URPN.	USPAT	OR	OFF	2006/02/10 12:58

3/3/06 10:53:25 AM

Page 1

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S29	0	("2004/0260976").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S30	0	("2004/0260972").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S31	0	("2004/0250032").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S32	0	("2004/0250031").URPN.	USPAT	OR	OFF	2006/02/10 12:59
S33	0	("2004/0250030").URPN.	USPAT	OR	OFF	2006/02/10 13:02
S34	0	("2004/0250029").URPN.	USPAT	OR	OFF	2006/02/10 13:11
S35	10	((batch\$4) with adjust\$4 with size) same (copy\$5 or redundan\$4 or mirror\$4)	US-PGPUB; USPAT; DERWENT	OR	ON	2006/02/10 13:12
S36	2	"6728484".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:04
S37	2	"6728848".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:05
S38	3	"6606694".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S39	0	os200300225760	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S40	0	us200300225760	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S41	0	"200300225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S42	0	"2003/00225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:06
S43	0	"US2003/00225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S44	0	"US2003/0225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S45	0	"US20030225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S46	2	"20030225760"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:07
S47	2	"20030135478"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:08

EAST Search History

S48	2	"20040107226"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S49	2	"20020059487"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S50	10	((("6732124") or ("5706470") or ("5140592") or ("6591351") or ("6148383"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/10 14:09
S51	1673	714/6.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/13 14:23
S52	1332	711/162.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/14 17:28
S53	1033	711/162.ccls. and (not 714/6.ccls.)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/15 14:34
S54	373	718/101.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/16 10:59
S55	423	709/233.ccls.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/16 11:00
S56	3	"20040205382"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 10:58
S57	0	("2004/0205382").URPN.	USPAT	OR	OFF	2006/02/22 10:58
S58	0	("2004/0205382").URPN.	USPAT	OR	OFF	2006/02/22 11:12
S59	20	((("5404508") or ("5596707") or ("5758359") or ("5796934") or ("5930824") or ("6247141") or ("6785786") or ("20030084372") or ("20030225646") or ("20040078628"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 11:24
S60	16	((("5513314") or ("6266784") or ("6728898") or ("5966730") or ("6163856") or ("6182086") or ("6636988") or ("20010056554"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 11:25
S61	2	"20030147344"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/22 13:00
S62	4	((("20030147344") or ("20030088698"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 12:46

EAST Search History

S63	22	((("4750106") or ("5325528") or ("5504861") or ("6052797") or ("6199074") or ("6282635") or ("6327671") or ("6442709") or ("6499112") or ("6553509") or ("20030041280"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 12:47
S64	4	((("6467024") or ("5720029"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 13:49
S65	4	((("5222219") or ("6493809") or ("6502205") or ("5900020"))).PN.	USPAT	OR	OFF	2006/02/23 13:51
S66	8	((("5051887") or ("5592618") or ("6618794") or ("6671705") or ("20010008008") or ("20030145179") or ("20030182525") or ("20040128442"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 13:52
S67	24	((("6816951") or ("6694447") or ("6691245") or ("6681339") or ("5615329") or ("6549921") or ("5928367") or ("20050071708") or ("6862671") or ("6728898") or ("6052797") or ("20040193952"))).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 14:01
S68	12	((("6816951") or ("6694447") or ("6691245") or ("6681339") or ("5615329") or ("6549921") or ("5928367") or ("20050071708") or ("6862671") or ("6728898") or ("6052797") or ("20040193952"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 14:13
S69	3	((("20010049749") or ("20030051111") or ("6591351"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 14:14
S70	8	((("6502205") or ("6931576") or ("20030101321") or ("6405294") or ("6397307") or ("6260125") or ("5592618") or ("6496908"))).PN.	US-PGPUB; USPAT	OR	OFF	2006/02/23 15:10
S71	0	("20330135478").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:17
S72	2	("20030135478").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:19
S73	2	("20020083120").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:19

EAST Search History

S74	5	((("6728848") or ("6606694")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:20
S75	7	((("6391576") or ("6173377") or ("6061770")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:21
S76	14	((("20040107226") or ("20020059487") or ("6732124") or ("5706470") or ("5140592") or ("6591351") or ("6148383")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:22
S77	10	((("20040107226") or ("20020059487") or ("6732124") or ("5706470") or ("5140592")).PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/23 15:25
S78	20	("4534030" "5371878" "5544347" "5555371" "5742792" "5771344" "5799141" "5889935" "5901327" "5903898" "5909548" "5933653" "6044444" "6052797" "6105029").PN. OR ("6442706").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/02/24 09:36
S79	74	"5140592"	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/24 11:09
S80	2	"5140592".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/24 11:09
S81	0	"08966633".ap.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/28 15:04
S82	6	"966633".ap.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/02/28 15:04
S83	2	("6148383").PN.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:48
S84	62	(failover and failback) with (memory or storage)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:53
S85	4	(failover and failback) with (memory or storage) and \$3commit\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 13:54
S86	4	(failover and failback) with (memory or storage) and \$3commit\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:02
S87	7	(failover) with (memory or storage) and \$3commit\$5 and failback	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:03

EAST Search History

S88	136	714/6.ccls. and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:19
S89	15	714/6.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:36
S90	0	("2005/0154937").URPN.	USPAT	OR	OFF	2006/03/02 14:04
S91	12	714/5.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:11
S92	0	("2005/0071708").URPN.	USPAT	OR	OFF	2006/03/02 14:12
S93	15	711/162.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:19
S94	0	718/101.ccls. and (failover or fail-over) and (fallback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 14:19
S95	1	("2002/0099916").URPN.	USPAT	OR	OFF	2006/03/02 14:21
S96	17	("5712970" "5724501" "5761705" "5895485" "5896492" "6199105" "6247099" "6279078" "6490659" "6578158").PN. OR ("6681339").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 14:27
S97	4	((("6816951") or ("6691245") or ("5615329") or ("5928367"))).PN.	USPAT	OR	OFF	2006/03/02 14:27
S98	14	("5555371" "5592618" "6078932" "6212531" "6223304" "6292905" "6363497" "6408310" "6516394" "6549921").PN. OR ("6694447").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 14:50
S99	2	"5978565".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:30
S10 0	15	failover with commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 1	21	(fail-over or failover) with commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 2	23	(fail-over or failover) with \$3commit\$3	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 3	30	(fail-over or failover) with \$3commit\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:31
S10 4	0	("2004/0268177").URPN.	USPAT	OR	OFF	2006/03/02 15:32

3/3/06 10:53:25 AM

Page 6

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S10 5	0	("2004/0260977").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 6	0	("2004/0260976").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 7	0	("2004/0260972").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 8	0	("2004/0250032").URPN.	USPAT	OR	OFF	2006/03/02 15:32
S10 9	0	("2004/0250030").URPN.	USPAT	OR	OFF	2006/03/02 15:33
S11 0	0	("2004/0250029").URPN.	USPAT	OR	OFF	2006/03/02 15:33
S11 1	31	714/6.ccls. and (failover or fail-over) and commit\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 15:57
S11 2	20	714/6.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:02
S11 3	8	714/5.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:03
S11 4	18	711/162.ccls. and (failover or fail-over) and commit\$4 and (primary with secondary)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 16:21
S11 5	7	714/6.ccls. and remirror	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:06
S11 6	82	714/6.ccls. and mirror\$4 and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:02
S11 7	146	master and slave and mirror\$4 and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:02
S11 8	129	(failover or fail-over) and (commit\$4 with (copy\$4 or replicat\$6 or mirror\$4))	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:07
S11 9	2	"6970987".pn.	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 17:43
S12 0	9	("20020010542" "20020099727" "20030105912" "5321703" "5675802" "5729733" "6334168" "6336173" "6502205").PN. OR ("6970987").URPN.	US-PGPUB; USPAT; USOCR	OR	OFF	2006/03/02 17:43
S12 1	29048	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:20

3/3/06 10:53:25 AM

Page 7

C:\Documents and Settings\CMcCarthy\My Documents\EAST\Workspaces\10456029.wsp

EAST Search History

S12 2	1603	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:20
S12 3	71	(primary or master) with (slave or secondary or backup) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory) and (\$3commit\$6)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:33
S12 4	2442	(prior or original) adj2 (primary or master or first) with (storage or memory)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:28
S12 5	12	((prior or original) adj2 (primary or master or first) with (storage or memory)) same ((slave or second\$4) with (memory or storage)) and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:29
S12 6	0	("2005/0015663").URPN.	USPAT	OR	OFF	2006/03/02 19:32
S12 7	1	"5537533".pn.	USPAT	OR	OFF	2006/03/02 19:32
S12 8	22	(local) with (remote) with (switch\$4 or revers\$4 or swap\$4) with (storage or memory) and (\$3commit\$6)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:33
S12 9	0	("2005/0198454").URPN.	USPAT	OR	OFF	2006/03/02 19:33
S13 0	75	"714"/\$3.ccls. and (failover or fail-over) and (failback or fallback or fail-back)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:41
S13 1	51	707/202.ccls. and (failover or fail-over)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:45
S13 2	3928	(failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:46
S13 3	759	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:47
S13 4	724	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5) and \$3commit\$6 and (memory or storage)	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:47
S13 5	759	((failover or fail-over) or (failback or fail-back or fallback) with \$3synchroniz\$5 with (memory or storage)) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48

EAST Search History

S13 6	755	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 7	0	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6 and mirror44	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 8	343	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (memory or storage)) and \$3commit\$6 and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:48
S13 9	343	((failover or fail-over) or (failback or fail-back or fallback) with resynchroniz\$5 with (second\$4 or remote or slave) with (memory or storage)) and \$3commit\$6 and mirror\$4	US-PGPUB; USPAT; DERWENT	OR	OFF	2006/03/02 19:49



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	PATENT NUMBER	GROUP ART UNIT	FILE WRAPPER LOCATION
10/456,053		2113	0120

Change of Address/Power of Attorney

The following fields have been set to Customer Number 22879 on

- Correspondence Address
- Power of Attorney

The address of record for Customer Number 22879 is:

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

The Practitioners of record for Customer Number 22879 are:

PTO INSTRUCTIONS:

Please take the following action when the correspondence address has been changed to a customer number:

- 1) Add 'ADDRESS CHANGE TO CUSTOMER NUMBER' on the next available content line of the File Jacket.
- 2) Put a line through the old address on the File Jacket and enter the Customer Number as the new address.
- 3) File this Notice in the File Jacket.

Please take the following action when the correspondence address has NOT been changed:

- 1) File this Notice in the File Jacket



0120

PTO/SB/21 (08-03)

Approved for use through 08/30/2003. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FORM (to be used for all correspondence after initial filing)	Application Number	10/456,053
	Filing Date	June 06, 2003
	First Named Inventor	Ji et al.
	Art Unit	
	Examiner Name	
Total Number of Pages in This Submission	Attorney Docket Number	200309044-1

ENCLOSURES (Check all that apply)		
☐ Fee Transmittal Form	☐ Drawing(s)	☐ After Allowance communication to Technology Center (TC)
☐ Fee Attached	☐ Licensing-related Papers	☐ Appeal Communication to Board of Appeals and Interferences
☐ Amendment/Reply	☐ Petition	☐ Appeal Communication to TC (Appeal Notice, Brief, Reply Brief)
☐ After Final	☐ Petition to Convert to a Provisional Application	☐ Proprietary Information
☐ Affidavits/declaration(s)	☐ Power of Attorney, Revocation	☐ Status Letter
☐ Extension of Time Request	☐ Change of Correspondence Address	☐ Other Enclosure(s) (please identify below):
☐ Express Abandonment Request	☐ Terminal Disclaimer	13 References and Two Return Postcards
☒ Information Disclosure Statement	☐ Request for Refund	
☐ Certified Copy of Priority Document(s)	☐ CD, Number of CD(s) _____	
☐ Response to Missing Parts/Incomplete Application	Remarks	
☐ Response to Missing Parts under 37 CFR 1.52 or 1.53		

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT	
Firm or Individual name	Derek J. Westberg (Reg. No. 40,872)
Signature	
Date	Nov. 18, 2003

CERTIFICATE OF TRANSMISSION/MAILING	
I hereby certify that this correspondence is being facsimile transmitted to the USPTO or deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the date shown below.	
Typed or printed name	Daniel Ratoff
Signature	
Date	Nov. 18, 2003


This collection of information is required by 37 CFR 1.5. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

PATENT APPLICATION

ATTORNEY DOCKET NO. 200309044-1



IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Ji et al.

Confirmation No.:

Application No.: 10/456,053

Examiner:

Filing Date: June 06, 2003

Group Art Unit:

Title: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

COMMISSIONER FOR PATENTS
Alexandria, VA 22313-1450

INFORMATION DISCLOSURE STATEMENT

Sir:

This Information Disclosure Statement is submitted:

- ☒ under 37 CFR 1.97(b), or
(Within three months of filing national application; or date of entry of national application; or before mailing date of first office action on the merits; whichever occurs last)
- ☐ under 37 CFR 1.97 (c) together with either a:
☐ Statement under 37 CFR 1.97(e), or
☐ a \$180.00 fee under 37 CFR 1.17(p), or
(After the CFR 1.97 (b) time period, but before final action or notice of allowance, whichever occurs first)
- ☐ under 37 CFR 1.97 (d) together with a:
☐ Statement under 37 CFR 1.97(e)(1) or (2), and
☐ a \$180.00 fee set forth in 37 CFR 1.17(p).
(Filed after final action, a notice of allowance, on or before before payment of the issue fee)

Please charge to Deposit Account **08-2025** the sum of \$0.00. At any time during the pendency of this application, please charge any fees required or credit any overpayment to Deposit Account **08-2025** pursuant to 37 CFR 1.25.

☒ Applicant(s) submit herewith Form PTO 1449 - Information Disclosure Citation together with copies, of patents, publications or other information of which applicant(s) are aware, which applicant(s) believe(s) may be material to the examination of this application and for which there may be a duty to disclose in accordance with 37 CFR 1.56.

☐ A concise explanation of the relevance of foreign language patents, foreign language publications and other foreign language information listed on PTO Form 1449, as presently understood by the individual(s) designated in 37 CFR 1.56 (c) most knowledgeable about the content is given on the attached sheet, or where a foreign language patent is cited in a search report or other action by a foreign patent office in a counterpart foreign application, an English language version of the search report or action which indicates the degree of relevance found by the foreign office is listed on form PTO 1449 and is enclosed herewith.

It is requested that the information disclosed herein be made of record in this application.

- ☒ I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to:
Commissioner for Patents, Alexandria, VA 22313-1450.
Date of Deposit: Nov. 18, 2003


- ☐ I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number _____ on _____
Number of pages: _____

Typed Name: Daniel Ratoff

Signature: 

Respectfully submitted,

Ji et al.

By 

Derek J. Westberg

Attorney/Agent for Applicant(s)

Reg. No. **40,872**

Date: **Nov. 18, 2003**

Telephone No.: **(408) 293-9000**



PATENT APPLICATION

Sheet 1 of 3

FORM PTO-1449

LIST OF PATENTS AND PUBLICATIONS FOR
APPLICANT'S INFORMATION DISCLOSURE
STATEMENT

(Use several sheets if necessary)

ATTY. DOCKET NO.

200309044-1

APPLICATION NO.

10/456,053

CONFIRMATION NO.

APPLICANT

Ji et al.

FILING DATE

06/06/03

GROUP

REFERENCE DESIGNATION

U.S. PATENT DOCUMENTS

EXAMINER INITIAL		DOCUMENT NUMBER	PUBLICATION DATE	NAME	Pages, Columns, Lines Where Relevant Passages or Figures Appear
	1A	5,544,347	Aug. 6, 1996	Yanai et al.	
	1B	5,742,792	Apr. 21, 1998	Yanai et al.	
	1C	5,909,692	Jun. 1, 1999	Yanai et al.	
	1D	6,092,066	Jul. 18, 2000	Ofek	
	1E	6,101,497	Aug. 8, 2000	Ofek	
	1F	6,108,748	Aug. 22, 2000	Ofek et al.	
	1G				
	1H				
	1I				
	1J				
	1K				

FOREIGN PATENT DOCUMENTS

		DOCUMENT NUMBER	PUBLICATION DATE	NAME OF PATENTEE OR APPLICANT	Pages/Columns/Lines Where Relevant Passages/Figures Appear	Check If Translation attached
	1L					
	1M					
	1N					
	1O					
	1P					

OTHER REFERENCES (including Author, Title, Date, Pertinent Pages, etc.)

1Q	Fay Chang, Minwen Ji, Shun-Tak A. Leung, John MacCormick, Sharon E. Perl and Li Zhang, Myriad: Cost-effective Disaster Tolerance, Proceedings of the FAST 2002 Conference on File and Storage Technologies, Monterey, CA, pp. 103-116, The USENIX Association, Berkeley, CA, January 2002.
1R	Chia Chao, Robert English, David Jacobson, Alexander Stepanov, and John Wilkes, Mime: a high performance parallel storage device with strong recovery guarantees, HP Laboratories Technical Report HPL-CSP-92-9 rev 1, 18 March 1992, revised 6 November 1992, Hewlett-Packard Company, Palo Alto, CA, 1992.
1S	Remote Copy Administrator's Guide and Reference, IBM DFSMS/MVS Version 1, Fourth Edition, Pages i to xv and pages 1 thru 170, International Business Machines Corporation, Raleigh, NC, December 1997.

EXAMINER

DATE CONSIDERED



PATENT APPLICATION

Sheet 2 of 3

FORM PTO-1449 LIST OF PATENTS AND PUBLICATIONS FOR APPLICANT'S INFORMATION DISCLOSURE STATEMENT (Use several sheets if necessary)	ATTY. DOCKET NO. 200309044-1	APPLICATION NO. 10/456,053	CONFIRMATION NO.
	APPLICANT Ji et al.		
	FILING DATE 06/06/03	GROUP	

REFERENCE DESIGNATION

U.S. PATENT DOCUMENTS

EXAMINER INITIAL		DOCUMENT NUMBER	PUBLICATION DATE	NAME	Pages, Columns, Lines Where Relevant Passages or Figures Appear
	2A				
	2B				
	2C				
	2D				
	2E				
	2F				
	2G				
	2H				
	2I				
	2J				
	2K				

FOREIGN PATENT DOCUMENTS

		DOCUMENT NUMBER	PUBLICATION DATE	NAME OF PATENTEE OR APPLICANT	Pages/Columns/Lines Where Relevant Passages/Figures Appear	Check if Translation attached
	2L					
	2M					
	2N					
	2O					
	2P					

OTHER REFERENCES (including Author, Title, Date, Pertinent Pages, etc.)


2Q	Hugo Patterson, Stephen Manley, Mike Federwisch, Dave Hitz, Steve Kleiman and Shane Owara, SnapMirror: File System Based Asynchronous Mirroring for Disaster Recovery, Proceedings of the FAST 2002 Conference on File and Storage Technologies, Monterey, CA, pp. 117-129, The USENIX Association, Berkeley, CA, January 2002.
2R	Symmetrix Remote Data Facility (SRDF) Product Description Guide, EMC Corporation, Hopkinton, MA, USA, 2000.
2S	Michael Stonebraker and Gerhard A. Schloss, Distributed Raid -- A New Multiple Copy Algorithm, UCB Electronics Research Laboratory Technical Report ERL-M89-56, Electronics Research Laboratory, University of California, Berkeley, CA, 1989.

EXAMINER	DATE CONSIDERED
----------	-----------------

Rev 02/03 (PTO1449)

PATENT APPLICATION

Sheet 3 of 3



FORM PTO-1449

LIST OF PATENTS AND PUBLICATIONS FOR
APPLICANT'S INFORMATION DISCLOSURE
STATEMENT

(Use several sheets if necessary)

ATTY. DOCKET NO.

200309044-1

APPLICATION NO.

10/456,053

CONFIRMATION NO.

APPLICANT

Ji et al.

FILING DATE

06/06/03

GROUP

REFERENCE DESIGNATION

U.S. PATENT DOCUMENTS

EXAMINER INITIAL		DOCUMENT NUMBER	PUBLICATION DATE	NAME	Pages, Columns, Lines Where Relevant Passages or Figures Appear
	3A				
	3B				
	3C				
	3D				
	3E				
	3F				
	3G				
	3H				
	3I				
	3J				
	3K				

FOREIGN PATENT DOCUMENTS

		DOCUMENT NUMBER	PUBLICATION DATE	NAME OF PATENTEE OR APPLICANT	Pages/Columns/Lines Where Relevant Passages/Figures Appear	Check if Translation attached
	3L					
	3M					
	3N					
	3O					
	3P					

OTHER REFERENCES (including Author, Title, Date, Pertinent Pages, etc.)

	3Q	EMC TimeFinder Product Description Guide, EMC Corporation, Hopkinton, Massachusetts, USA, 1998.
	3R	
	3S	

EXAMINER

DATE CONSIDERED

Rev 02/03 (PTO1449)

PATENT APPLICATION

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION

ATTORNEY DOCKET NO. 200309044-1

As below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

the specification of which is attached hereto unless the following box is checked:

() was filed on _____ as US Application No. or PCT International Application
Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
			YES: _____ NO: _____
			YES: _____ NO: _____

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION NUMBER	FILING DATE

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879

Place Customer
Number Bar Code
Label hereSend Correspondence to:
HEWLETT-PACKARD DEVELOPMENT COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Direct Telephone Calls To:

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Minwen Ji Citizenship: CanadianResidence: Sunnyvale, CAPost Office Address: 862 Hydrangea Ct., Sunnyval, CA 94086-8115Inventor's Signature [Signature] Date 6/6/03

Rev 02/03 (DocPwr)

(Use Page Two For Additional Inventor(s) Signature(s))

Page 1 of 2

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)

ATTORNEY DOCKET NO. 200309044-1

Full Name of # 2 joint inventor: Alistair Veitch Citizenship: New Zealand
Residence: Mountain View, CA
Post Office Address: 2420 Whitney Dr., Mountain View, CA 94043-4133
Inventor's Signature: A. C. Veitch Date: June 6, 2003

Full Name of # 3 joint inventor: John Wilkes Citizenship: United Kingdom
Residence: Palo Alto, CA
Post Office Address: 3788 Park Blvd., Palo Alto, CA 94306-3430
Inventor's Signature: _____ Date: _____

Full Name of # 4 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 5 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 6 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 7 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 8 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)

ATTORNEY DOCKET NO. 200309044-1

Full Name of # 2 joint inventor: Alistair Veitch Citizenship: New Zealand
Residence: Mountain View, CA
Post Office Address: 2420 Whitney Dr., Mountain View, CA 94043-4133

Inventor's Signature _____ Date _____

Full Name of # 3 joint inventor: John Wilkes Citizenship: United Kingdom
Residence: Palo Alto, CA
Post Office Address: 3788 Park Blvd., Palo Alto, CA 94306-3430

John Wilkes 10 June 2003
Inventor's Signature _____ Date _____

Full Name of # 4 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 5 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 6 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 7 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 8 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____

Inventor's Signature _____ Date _____



UNITED STATES PATENT AND TRADEMARK OFFICE


UNITED STATES DEPARTMENT OF COMMERCE
 United States Patent and Trademark Office
 Address: COMMISSIONER OF PATENTS AND TRADEMARKS
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
 www.uspto.gov

APPLICATION NUMBER	FILING OR 371 (c) DATE	FIRST NAMED APPLICANT	ATTORNEY DOCKET NUMBER
10/456,053	06/06/2003	Minwen Ji	200309044-1

HEWLETT-PACKARD DEVELOPMENT COMPANY
 Intellectual Property Administration
 P.O. Box 272400
 Fort Collins, CO 80527-2400

CONFIRMATION NO. 5067

FORMALITIES LETTER



OC000000010590314

Date Mailed: 07/29/2003

NOTICE TO FILE MISSING PARTS OF NONPROVISIONAL APPLICATION

FILED UNDER 37 CFR 1.53(b)

*Filing Date Granted*Items Required To Avoid Abandonment:

An application number and filing date have been accorded to this application. The item(s) indicated below, however, are missing. Applicant is given **TWO MONTHS** from the date of this Notice within which to file all required items and pay any fees required below to avoid abandonment. Extensions of time may be obtained by filing a petition accompanied by the extension fee under the provisions of 37 CFR 1.136(a).

- The signature of the following inventor(s) is missing from the oath or declaration:
John Wilkes
- To avoid abandonment, a late filing fee or oath or declaration surcharge as set forth in 37 CFR 1.16(e) of \$130 for a non-small entity, must be submitted with the missing items identified in this letter.

SUMMARY OF FEES DUE:

Total additional fee(s) required for this application is **\$130** for a Large Entity

- **\$130** Late oath or declaration Surcharge.

*A copy of this notice **MUST** be returned with the reply.*

Customer Service Center
 Initial Patent Examination Division (703) 308-1202

PART 3 - OFFICE COPY

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P. O. Box 272400
Fort Collins, Colorado 80527-2400

06-0910356053.060603 A

PATENT APPLICATION

ATTORNEY DOCKET NO. 200309044-1

01919 U.S. PTO



IN THE U.S. PATENT AND TRADEMARK OFFICE
Patent Application Transmittal Letter

06/06/03

COMMISSIONER FOR PATENTS
Alexandria, VA 22313-1450

Sir:

Transmitted herewith for filing under 37 CFR 1.53(b) is a(n): ☒ Utility ☐ Design

☒ original patent application,

☐ continuation-in-part application



INVENTOR(S): Ji et al.

TITLE: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

Enclosed are:

- ☒ The Declaration and Power of Attorney. ☐ signed ☒ unsigned or partially signed
☒ 14 sheets of drawings (one set) ☐ Associate Power of Attorney
☐ Form PTO-1449 ☐ Information Disclosure Statement and Form PTO-1449
☐ Priority document(s) ☐ (Other) (fee \$)

CLAIMS AS FILED BY OTHER THAN A SMALL ENTITY				
(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) TOTALS
TOTAL CLAIMS	47 — 20	27	X \$18	\$ 486
INDEPENDENT CLAIMS	2 — 3	0	X \$84	\$ 0
ANY MULTIPLE DEPENDENT CLAIMS	0		\$280	\$ 0
BASIC FEE: Design (\$330.00); Utility (\$750.00)				\$ 750
TOTAL FILING FEE				\$ 1,236
OTHER FEES				\$
TOTAL CHARGES TO DEPOSIT ACCOUNT				\$ 1,236

Charge \$ 1,236 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

"Express Mail" label no. EU687335939US

Date of Deposit June 6, 2003

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

By

Typed Name: Daniel Ratoff

Respectfully submitted,

Ji et al

By

Derek J. Westberg

Attorney/Agent for Applicant(s)

Reg. No. 40,872

Date: June 6, 2003

Telephone No.: (408) 293-9090

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P. O. Box 272400
Fort Collins, Colorado 80527-2400

06-0910356053.060603 A

PATENT APPLICATION

ATTORNEY DOCKET NO. 200309044-1

01919 U.S. PTO



IN THE U.S. PATENT AND TRADEMARK OFFICE
Patent Application Transmittal Letter

06/06/03

COMMISSIONER FOR PATENTS
Alexandria, VA 22313-1450

Sir:

Transmitted herewith for filing under 37 CFR 1.53(b) is a(n): ☒ Utility ☐ Design

☒ original patent application,

☐ continuation-in-part application



INVENTOR(S): Ji et al.

TITLE: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

Enclosed are:

- ☒ The Declaration and Power of Attorney. ☐ signed ☒ unsigned or partially signed
☒ 14 sheets of drawings (one set) ☐ Associate Power of Attorney
☐ Form PTO-1449 ☐ Information Disclosure Statement and Form PTO-1449
☐ Priority document(s) ☐ (Other) (fee \$)

CLAIMS AS FILED BY OTHER THAN A SMALL ENTITY				
(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) TOTALS
TOTAL CLAIMS	47 — 20	27	X \$18	\$ 486
INDEPENDENT CLAIMS	2 — 3	0	X \$84	\$ 0
ANY MULTIPLE DEPENDENT CLAIMS	0		\$280	\$ 0
BASIC FEE: Design (\$330.00); Utility (\$750.00)				\$ 750
TOTAL FILING FEE				\$ 1,236
OTHER FEES				\$
TOTAL CHARGES TO DEPOSIT ACCOUNT				\$ 1,236

Charge \$ 1,236 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

"Express Mail" label no. EU687335939US

Date of Deposit June 6, 2003

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

By

Typed Name: Daniel Ratoff

Respectfully submitted,

Ji et al

By

Derek J. Westberg

Attorney/Agent for Applicant(s)

Reg. No. 40,872

Date: June 6, 2003

Telephone No.: (408) 293-9090

FAULT-TOLERANT DATA REDUNDANCY TECHNIQUERelated Applications:

The following applications disclose related subject matter: U.S. Application No. (Attorney Docket No. 100204276-1), filed (on the same day as this application) and entitled, "Asynchronous Data Redundancy Technique"; U.S. Application No. (Attorney Docket No. 200309042-1), filed (on the same day as this application) and entitled, "Redundant Data Consistency After Failover"; U.S. Application No. (Attorney Docket No. 200309043-1), filed (on the same day as this application) and entitled, "Distributed Data Redundancy Operations"; U.S. Application No. (Attorney Docket No. 200309045-1), filed (on the same day as this application) and entitled, "Adaptive Batch Sizing for Asynchronous Data Redundancy"; U.S. Application No. (Attorney Docket No. 200309046-1), filed (on the same day as this application) and entitled, "State Machine and System for Data Redundancy"; U.S. Application No. (Attorney Docket No. 200309047-1), filed (on the same day as this application) and entitled, "Batched, Asynchronous Data Redundancy Technique"; U.S. Application No. (Attorney Docket No. 200309499-1), filed (on the same day as this application) and entitled, "Data Redundancy Using Portal and Host Computer"; the contents of all of which are hereby incorporated by reference.

Background of the Invention:

The present invention relates to the field of data storage. More particularly, the present invention relates to techniques for redundant data storage.

Remote mirroring is a data redundancy technique for coping with storage system failures. A copy of data, sometimes referred to as a 'primary' or 'local' copy, is updated, for example, as it is accessed by an application program. A redundant copy of the data, sometimes referred to as a 'secondary' or 'slave' copy of the data, usually at a remote site, is updated as well. When a failure occurs that renders the primary copy unusable or inaccessible, the data can be restored from the secondary copy, or accessed directly from there.

Conventional techniques for remote mirroring tend to maintain the primary and secondary copies of the data synchronized. However, such techniques do not cope well with unexpected circumstances such as lengthy communication delays to the remote site, buffers filled to capacity, failures, and so forth.

Therefore, what is needed is an improved technique for redundant data storage. It is to this end that the present invention is directed.

Summary of the Invention:

5 The invention provides methods and apparatus for performing data redundancy operations in a fault-tolerant manner. In one aspect, a primary data storage facility stores a primary copy of data and a secondary facility stores data that is redundant of the primary copy of the data. The primary facility includes a first redundancy appliance that receives a sequence of write requests and stores data for
10 the sequence of write requests in storage associated with the primary storage facility. A second redundancy appliance shadows the first redundancy appliance and assumes the role of the first redundancy appliance in the event of a fault at the first redundancy appliance. In this way, fault tolerance is increased by the presence of the second, shadow appliance.

15 The second redundancy appliance may maintain a shadow copy of data stored by the first redundancy appliance. In which case, the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance. Alternately, the second redundancy appliance may access the data stored by the first redundancy appliance when the
20 second redundancy appliance assumes the role of the first redundancy appliance.

These and other aspects of the invention are explained in more detail herein.

Brief Description of the Drawings:

Figure 1 illustrates a computer system including a primary data storage facility
25 and a secondary data storage facility in which the present invention may be implemented;

Figure 2 illustrates operation of the primary and secondary storage facility of Figure 1 in accordance with an embodiment of the present invention;

Figure 3 illustrates the computer system of Figure 1 in more detail including
30 write queues at the primary and secondary data storage facilities in accordance with an embodiment of the present invention;

Figure 4 illustrates an exemplary relationship between communication bandwidth and batch size that may be utilized in accordance with an embodiment of the present invention;

Figures 5A-B illustrate send and receive barriers in accordance with an embodiment of the present invention;

Figure 6 illustrates a flow diagram for queuing and applying a batch of transactions at the secondary facility in accordance with an embodiment of the present invention;

Figure 7 illustrates a state machine for controlling operation of the primary and/or secondary data storage facilities of Figure 1 in accordance with an embodiment of the present invention;

Figure 8 illustrates an example of update and back-up copy propagation during failover and recovery in accordance with an embodiment of the present invention;

Figure 9 illustrates a second example of update and back-up copy propagation example during failover and recovery in accordance with an embodiment of the present invention;

Figure 10 illustrates a third example of update and back-up copy propagation example during failover and recovery in accordance with an embodiment of the present invention;

Figure 11 illustrates primary and secondary storage facilities in which redundant elements are provided in accordance with an embodiment of the present invention;

Figure 12 illustrates primary and secondary storage facilities in which data storage is distributed in accordance with an embodiment of the present invention;

Figure 13 illustrates primary and secondary storage facilities including network portal redundancy appliances in accordance with an embodiment of the present invention; and

Figure 14 illustrates a flow diagram of a method for testing a data redundancy system in accordance with an embodiment of the present invention.

Detailed Description of a Preferred Embodiment:

The invention provides methods and apparatus for performing data redundancy operations in a fault-tolerant manner. A primary data storage facility stores a primary copy of data and a secondary facility stores data that is redundant of the primary copy of the data. The primary facility includes a first redundancy appliance that receives a sequence of write requests and stores data for the sequence of write requests in storage associated with the primary storage facility. A second

redundancy appliance shadows the first redundancy appliance and assumes the role of the first redundancy appliance in the event of a fault at the first redundancy appliance. In this way, fault tolerance is increased by the presence of the second, shadow appliance.

5 The invention can be applied to any computer system in which a primary copy of data is backed up by data that is redundant of the primary copy. For example, the primary copy may be stored at a primary data storage facility, while redundant data may be stored at one or more secondary storage facilities. The data storage facilities can include any type of data storage, such as volatile or non-volatile memory,
10 including random access memory, flash memory, magnetic tape or disk, an array of disk drives and so forth. The primary and secondary storage facilities are positioned at different locations, which are generally remote from one another. Thus, the storage facilities communicate via a network or via a direct communication link. Exemplary communication networks include: local area networks (LANs), metropolitan area
15 networks (MANs), wide area networks (WANs), storage area networks (SANs), the Internet and so forth.

Figure 1 illustrates a computer system 100 by which the present invention may be implemented. The system 100 includes a primary data storage facility 102, a secondary data storage facility 104 and a communication medium 106, such as a
20 network, for interconnecting the primary and secondary storage facilities 102 and 104.

Additional devices, such as one or more computer(s) 108 (e.g., a host computer, a workstation or a server), may communicate with the primary data storage facility 102 (e.g., via communication medium 110). While Figure 1 illustrates the communication medium 106 and the communication medium 110 as being separate,
25 they may be combined. For example, communication between the computer 108 and the primary facility 102 may be through the same network as is used for the primary storage facility 102 and secondary storage facility 104 to communicate.

One or more applications operating at the computer 108 may access the primary data storage facility 102 for performing write or read transactions to or from data objects, such as files or storage volumes, stored at the facility 102. More
30 particularly, the computer 108 may retrieve a copy of a data object by issuing a read request to the facility 102. Also, when a data object at the computer 108 is ready for storage at the facility 102, the computer 108 may issue a write request to the facility 102. For example, the computer 108 may request storage of a file undergoing

modification by the computer 108. While a single computer 108 is illustrated in Figure 1, it will be apparent that multiple computers may access the data storage facilities 102 and 104. In addition, a computer system 100 may include any number of devices that retrieve, modify and/or generate data and any number of primary and secondary storage facilities. Further, a device, such as a workstation or server, may also function as a storage facility. Still further, a storage facility may function as a primary storage facility for some data and as a secondary storage facility for other data, and a storage facility may function as a computer system, generating storage requests (e.g., as part of a backup process). The connections between the various components shown in Figure 1 are purely exemplary: any other topology, including direct connections, multiple networks, multiple network fabrics, etcetera, may be used.

For increasing data reliability in the event of a fault at the primary storage facility 102, data that is redundant of data stored at the primary facility 102 is stored at the secondary facility 104. For example, the secondary facility 104 may store a mirrored copy of the data. Alternately, the redundant data may be arranged according to a redundancy scheme in which redundant data is distributed among or striped across multiple storage devices or facilities. For example, the redundant data may be stored at the secondary facility 104 in accordance with Redundant Array of Inexpensive Disks (RAID) techniques, such as RAID levels 2, 3, 4 or 5. Further, one or more additional secondary storage facilities may be provided, in which each stores only a portion of the data stored at the primary 102 (thus, providing a distributed redundant copy) or where each stores a complete copy of the data (thus, providing multiple redundant copies).

In absence of a fault at the primary facility 102, the computer 108 generally does not direct write and read accesses to the secondary storage facility 104. Rather, for performing write and read operations, the computer 108 accesses the primary storage facility 102. The primary facility 102 and the secondary facility 104 then interact to provide redundant data at the secondary facility 104. In the event of a fault at the primary storage facility 102, lost data may then be reconstructed from the redundant data stored at the secondary facility 104 and delivered to the computer 108, or another computer (not shown) may be used to access data at the secondary facility 104 after failover.

Figure 2 illustrates operation of the primary and secondary storage facilities 102 and 104 of Figure 1 in accordance with an aspect of the present invention. A redundancy appliance 202 at the primary facility 102 is illustrated in Figure 2 along with a redundancy appliance 204 at the secondary facility 104. It will be apparent
5 that the appliances 202 and 204 may be implemented by (amongst other examples) appropriately configured hardware, software or firmware in disk arrays, storage devices, hosts (e.g., computer 108), in-host I/O bus adapters, network switches, network hubs, or combination thereof, which may be dedicated to perform the functions of the appliances 202 and 204 as described herein, or which may have
10 shared functionality.

As used herein, a “local” storage facility is typically physically positioned in proximity to the computer 108, whereas a “remote” storage facility is other than the local storage facility and is typically more distant from the computer 108. A “primary” storage facility is currently providing services with respect to a primary
15 copy of the data, while a “secondary” storage facility is other than the primary storage facility and typically acts as a backup by storing data redundantly. Under normal conditions, e.g., in the absence of a fault at the local facility, the local facility typically serves as the primary facility. However, in the event of a fault at the local facility (or under other conditions), the remote facility may assume the role of the primary
20 facility, as explained in more detail herein. Also, the remote facility may function as a primary facility for some data storage operations and as a secondary data storage facility for other data storage operations.

Referring to Figure 2, when a local facility also serves as the primary facility 102, a write request at the primary facility 102 (e.g., issued by the computer 108)
25 causes a write record to be written into a primary log 206 at the primary facility 102. The write-ordering of the requests in the primary log 206 may be preserved by writing the records synchronously (in the order of occurrence), or by other means, such as appropriate record-keeping. In addition, the corresponding data for the request is written to a primary copy of the data 208, which may be stored as one or more logical
30 units (LUs) at the primary facility 102. An acknowledgement may then be sent to the computer 108 indicating the request was successfully stored by the primary facility 102. In what follows, we use logical units (LUs) as exemplary; any convenient storage entity may be used, including other types of storage devices, files, and databases.

The write record is preferably written to the primary log 206 synchronously with the write request to the primary copy of the data 208 so as to preserve the write-ordering of the requests, however, the data may be written to the primary log 206 asynchronously. The primary log 206 may be stored, for example, in a dedicated
5 storage device (e.g., a disk drive, disk array or section of non-volatile memory (NVRAM)) associated with the appliance 202 at the primary facility 102 or in a storage device that is accessible via a Storage Area Network (SAN), and may be shared with other uses. Preferably, at least the tail portion (i.e., the most recently appended-to part) of the primary log 206 is stored in NVRAM; either because all of it
10 is, or because the log is stored on a device equipped with a non-volatile memory. Preferably, the log 206 is stored in a storage device that is disjoint from any device used to store the primary copy 208 of the data.

The secondary facility 104 may include a redundancy appliance 204, a transaction log 210 and a data repository, e.g., one or more LUs 212.

15 Figure 3 illustrates the primary and secondary storage facilities 102 and 104 of the computer system 100 of Figure 1 in more detail. As shown in Figure 3, the primary storage facility 102 includes a primary storage controller 112, a local mass-storage media 114 and a write transaction queue 116. The primary controller 112 includes a processor for controlling operations of the primary storage facility 102,
20 including the storage of data in the mass-storage media 114 and the forwarding of data to the secondary storage facility 104 and, thus, performs the functions of the appliance 202 (Figure 2). The storage media 114 generally stores the primary copy 208 (Figure 2) and may include, for example, a disk drive or disk array. The write queue 116 generally stores the primary log 206 (Figure 2) and may be stored in a disk
25 or disk array associated with the primary storage facility 102; preferably, the write queue 116 is equipped with a non-volatile RAM and is disjoint from the local mass-storage 114 which holds the primary data copy. The primary and secondary storage controllers may be replicated, distributed, mirrored, or otherwise constructed using any of the techniques known in the art for building storage systems.

30 As mentioned, to store data at the primary storage facility 102, write requests are issued to the primary facility 102. In response, the storage facility 102 stores the data in its local storage media 114. In addition, when the data is also to be stored redundantly at the second storage facility 104, write transactions for the data are

inserted into the write queue 116, where they are queued for communication to the secondary data storage facility 104 via communication medium 106 (Figure 1).

The write queue 116 may function as a first-in, first-out buffer (FIFO) for write transactions. In one embodiment, the write transactions are immediately
5 forwarded from the write queue 116 to the secondary facility 104. In this embodiment, the write transactions may be forwarded in the order they are received by the primary facility.

In another embodiment, a sequence of “snapshots” of the primary LU 208 may be implemented in the log 206. The snapshots may include only the changed data, or
10 they may include a complete copy of the data that is brought up to date when the snapshot is taken (typically – and preferably – by being a mirrored copy of the data that is kept almost up to date so that this does not take too long). Although the invention described herein is preferably implemented using a log, the snapshots can be implemented in another manner.

15 In another embodiment, the primary storage facility 102 delays forwarding write transactions to the secondary facility 104. In this embodiment, the write transactions are preferably grouped into send batches prior to forwarding them. Overwrites within a send batch may be permitted, though preferably not across batch boundaries. More particularly, a batch of write transactions may be collected over
20 successive time intervals. The batches are, thus, formed one after the other. For example, as shown in Figure 3, a batch n is formed, then a batch n+1, then a batch n+2, and so forth. Write transactions received during an interval are assigned to the corresponding send batch.

In one aspect, all of a send batch may be forwarded to the secondary storage
25 facility before any of a next send batch is forwarded. Further, the send batches may be forwarded in the order of their formation or in another order. Also, more than one send batch may be forwarded at any one time.

The size of the batches may be based on collection of a predetermined count or aggregate size of write transactions into each batch or a predetermined amount of
30 data to be transferred by the batch. Alternately, the size of the batches may be determined by the duration of successive time intervals over which the batches of write transactions are collected. For example, the intervals may be measured according to time-intervals, e.g., ten or thirty seconds, during which the transactions are to be collected.

If a write transaction received during the interval affects the same data as an earlier operation received during the same interval (and, thus, the later-received operation overwrites the prior data), the later-received operation may replace the earlier operation in the send batch. Multiple write transactions may affect the same data, for example, where the computer 108 issues write requests to store intermediate versions of a data object while the data object is undergoing revision by computer 108 (Figure 1).

By allowing overwrites at the primary facility 102, the communication bandwidth required between the primary and secondary facility 104 may be reduced because the replaced write transactions are not forwarded. However, collecting write transactions at the primary server 102 tends to increase the quantity of data that could be lost should a failure occur at the primary server 102. This is because write transactions queued at the primary facility 102 reflect changes to the data which have not yet been propagated to the secondary facility 104. Accordingly, write transactions not yet propagated to the secondary facility 104 may be lost in the event of a failure at the primary facility 102.

Accordingly, the size of send batches (and whether write transactions are to be queued at the primary server 102) may be determined based on bandwidth availability between the storage facilities 102 and 104 and/or on the potential adverse consequences of the loss of write transactions in the event of a failure. Further, the batch size may be adjusted adaptively, based on these same considerations.

In one aspect, the level of communication bandwidth available in the medium 106 (Figure 1) may be detected and used for determining the batch size, in which case, the size of the send batches may be based on a level of traffic detected on the medium 106. When the traffic is heavy, a larger batch size will tend to reduce the added burden on the medium 106. Thus, to conserve communication bandwidth by allowing more overwrites during times of heavy network traffic, the send batch sizes may be increased. Conversely, when the traffic is light, a smaller batch size may be accommodated. Thus, batch size may be reduced in times of lighter traffic. This scheme may be used, for example, where the communication medium 106 is shared by other entities.

In another aspect, the communication medium may be monitored to determine when traffic is sufficiently low that the batch can be accommodated immediately. For example, where the communication medium 106 includes a link dedicated to

communications between the first and second facilities, the link may be monitored to determine when it is available (e.g., when it becomes idle). Upon the link becoming available, the current batch may be completed and forwarded along the link.

In yet another aspect, the size of send batches may be based on the communication bandwidth consumed by forwarding the batches, in which case, the batch size may be adjusted so as to optimize the trade-off between batch size and communication bandwidth. As mentioned, a larger batch size tends to reduce the bandwidth required to forward the batch by increasing the number of overwrites that may occur, but also increases the amount of data that may potentially be lost if a failure prevents the batch from being forwarded to the secondary facility 104. Figure 4 illustrates an exemplary diagram showing a relationship between communication bandwidth and batch size that may be utilized. This relationship may be represented by a function and may be determined experimentally, for example, by measuring the bandwidth consumed for each of several different batch sizes. As shown in Figure 4, increasing the batch size may have a dramatic effect on reducing bandwidth, as shown by the steep slope in the graph, up to a certain point at which the slope is reduced (e.g., an inflection in the graph is reached). Beyond this point, further increases in batch size may have a diminished effect on bandwidth and, thus, the potential for loss of data in the event of a failure will likely tend to outweigh any additional bandwidth savings. A preferred batch size coincides with the change in slope or inflection.

In a further aspect, the send batch sizes may be selected based on the expected time between failures that inhibit forwarding of the send batches to the secondary storage facility 104. For example, the mean time between failures for the primary facility and/or the communication medium 106 may be determined (e.g., experimentally or based on manufacturer's data). Where the expected time between failures is relatively long, this indicates that failures will occur rarely. Thus, a larger batch size may be used since fewer batches will be lost due to such failures. However, where the expected time between failures is short, this indicates that such failures may occur frequently. Thus, a smaller batch size may be used since this data is subject to loss in the event of a failure. Further, once a batch size has been selected, it may be adjusted if further monitoring of the time between failures indicates that failures occur more or less frequently than originally anticipated. For example, where monitoring (e.g., by the primary controller 112) indicates that failures occur more frequently than previously expected, the batch size may be automatically reduced

(e.g., by the primary controller 112) and, where failures occur less frequently than previously expected, the batch size may be automatically increased.

When a send batch is completed, new write transactions are collected into the next send batch. For example, when the batch n is completed, subsequent write
5 transactions are collected into batch n+1. Also, once completed, the batch n is ready for forwarding to the secondary facility 104. Preferably, completed batches are forwarded as soon as practical so as to minimize data loss should a failure occur at the primary facility 102 before a batch is forwarded to the secondary facility 104. Accordingly, the batches are preferably communicated to the secondary facility 104 in
10 the order in which they are formed (i.e. n, n+1, n+2, n+3, etc.).

As is also shown in Figure 3, the secondary facility 104 includes a secondary controller 118, mass-storage media 120, which generally stores the redundant data 212 (Figure 2) and a write transaction queue 122, which generally stores the log 210 (Figure 2). Similarly to the primary storage facility 102, the controller 118 of the
15 secondary storage facility 104 includes a processor for controlling operations of the secondary storage facility 104 and, thus, performs the functions of the appliance 204 (Figure 2). This includes controlling the reception of transactions from the primary storage facility 102 and controlling the storage of data in the mass-storage media 120. The storage media 120 may include, for example, a hard disk array.

20 In response to receiving write transactions from the primary storage facility 102, the secondary storage facility 104 queues the operations in its write queue 122 and then stores the updated data in its storage media 120. However, the write transactions may not be applied to the redundant data (and, thus, remain in the queue 122) until after a delay has elapsed or a specified event has occurred (or until a
25 combination thereof occurs). Delaying application of the write transactions inhibits the propagation of errors to the redundant data. For example, a software error may occur at the primary facility 102 or at the computer 108 that results in sending corrupted data to the primary copy. By delaying application of the corrupted data to the redundant data at the secondary facility 104, propagation of the error may be
30 halted during the delay interval by avoiding applying the corrupted data.

The write transactions may be queued at the secondary facility 104 in the same order and form in which they are received from the primary facility 102. Thus, where the primary facility 102 forwards the write transactions one at a time, they may be queued individually at the secondary facility 104 in the order they are received.

Similarly, where the primary facility 102 forwards the write transactions in batches (e.g., n , $n+1$, $n+2$, etc.), the write transactions may be queued at the secondary facility 104 according to the same batches and in the order in which they are received.

5 In one aspect, the write transactions received from the primary facility 102 are collected into one or more receive batches of transactions at the secondary facility 104. The boundaries of the receive batches collected at the secondary facility need not bear a relationship to those of the send batches collected at the primary facility 102. The receive batches are shown in Figure 3 by the batches m , $m+1$, $m+2$, etc. Thus, where the write transactions are received one at a time, multiple operations may
10 be collected into a receive batch. Where the write transactions are received according to send batches (e.g., n , $n+1$, $n+2$, etc.) multiple send batches may be applied as a whole to the mass-storage media 120 (i.e. all of the transactions in that batch are applied or none are). Applying the write transactions as a whole may be performed, for example, by repeatedly re-applying a log of write transactions until all are applied,
15 storing data for the write transactions and a map of the data and then changing the map or by using copy-on-write techniques (in which a prior version of the data is saved in case it is needed again). Thus, overwrites may be allowed across receive batches where write transactions are replaced by later-received write transactions that affect the same data and the receive batches that contain such overwrites are
20 combined into a single receive batch, which will be applied as a whole. Applying the entire batch as a whole avoids the redundant data becoming internally inconsistent – and unrecoverable – as might otherwise occur if the ordering of the write transactions is not preserved across batch boundaries.

Figures 5A-B illustrate send and receive barriers in accordance with an aspect
25 of the present invention. Send barriers may be generated to indicate the boundaries of send batches. As mentioned, overwrites may be allowed within a batch, but not across batches. The send barrier of Figures 5A-B indicates the start of a send batch to which new write transactions are to be appended. As mentioned, the size of the send batches can be based on a number of criteria, such as the number of transactions, the
30 amount of data to transfer at a time or a time interval.

Receive barriers bound the sets of transactions or data blocks that are to be applied as a whole (i.e. all the transactions are applied or none are). A receive barrier may initially be associated with each write transaction; that is, each data block may be a receive batch by itself. When a block in the same send batch is overwritten, the

earlier write record for that transaction is removed from the queue 122 as are any
receive barriers for blocks written between the old copy and the new write transaction.
This merges the transactions for blocks that had been separated by receive barriers
into the same receive batch. Thus, depending on the circumstances, receive batches
5 may be smaller than send batches.

As shown in Figure 5A, a series of data blocks A, B, C and D are written to.
The corresponding transactions may be entered into the write queue 116 (Figure 2) in
the order in which the transactions occur. In the example, of Figure 5A, the order is
A-D-A-B-C-A, where the last transaction affecting data block A is shown being
10 appended to the queue 116. The send barrier indicates the end of the prior batch and
the start of the current batch. Also, shown in Figure 5A are receive barriers that may
be associated with the transactions. When the last transaction to data block A is
appended, the prior transaction within the same send batch may be removed (i.e.
overwritten). This is shown in Figure 5B, in which the prior transaction to block A
15 has been removed. In addition, Figure 5B illustrates that the receive barriers for
blocks occurring between the removed transaction and the new ("overwriting")
transaction are removed. As such, these blocks need to be written at the secondary
facility 104 as a whole to preserve the write-ordering of transactions. Thus, in the
example, the transactions to blocks B, C and A are to be written at the second facility
20 104 as a whole.

Receive batches may be merged at the secondary, by concatenating two or
more adjacent receive batches together, and eliminating data overwritten in a later
receive batch of those concatenated together. This may be used to reduce the amount
of space needed at the secondary; to exploit overwrite activity; to save on metadata
25 information; to reduce processing load; or for any other reason. Batch concatenation
may be triggered by detecting one or more of these conditions; such detection may
occur at the arrival of a new batch; periodically; on demand; or at any other
convenient or appropriate time.

The controller 112 preferably keeps track of the locations of the send barriers
30 and the receive barriers. So that the secondary facility 104 can identify transactions to
be applied as a whole, the controller 112 also forwards information sufficient to
enable the secondary facility 104 to identify the receive barriers. For example, this
information may be sent with the send batch, but may only be required if the send
batch and receive barriers do not coincide.

The delay associated with the write queue 122 at the secondary facility 104 may be determined in a number of different ways. For example, where the write transactions are received and applied individually, a timestamp may be associated with each transaction. The timestamp may be created when the transaction is queued at the primary facility 102 or when the transaction is received by the secondary facility 104. Each timestamp may indicate the then-current time, such as time of day. When a timestamp reaches a predetermined age, e.g., 30 seconds, 10 minutes, or 1 day, the timestamp expires, though not all timestamps need to expire after the same amount of time. For example, a timestamp may incorporate its own expiration time. When the timestamp expires, the redundant data 212 (Figure 2) may be updated in accordance with the transaction. Similarly, where write transactions are received and applied according to send batches (e.g., n, n+1, n+2, etc.) formed at the primary facility 102, a timestamp may be associated with each send batch. The timestamp may be created, for example, when the batch is formed at the primary facility 102 or when the batch is received at the secondary facility 104. Where a single timestamp is associated with multiple transactions, its precision can be approximate. For example, the timestamp may be created when a first, last or an intermediate transaction within the send batch is queued or communicated. Then, when the timestamp expires, the redundant data may be updated in accordance with the batch of operations, where each batch is applied as a whole.

Where the multiple operations are collected in receive batches, a timestamp may be associated with each receive batch. For example, the timestamp for a batch may be formed when the batch is completed. Then, when the timestamp expires (e.g., when it becomes 30 minutes old), the redundant data is updated in accordance with the batch of operations, where each batch is applied as a whole.

Rather than waiting to apply the write transactions to the redundant data according to elapsed time, the write transactions may be queued at the secondary facility 104 until a specified event occurs that indicates that the transactions are safe to apply. For example, a data integrity verification such as virus detection, intrusion detection, verifying a checksum or verification of network logs may be performed on the data to be updated or the original copy, or both, before the operations are applied to determine whether irregularities may indicate that the data may possibly be corrupted. These checks may be performed, for example, at the secondary facility 104 (e.g., by the controller 118) based on transactions in the queue 122 or at the

primary facility 102 (e.g., by the controller 112) based on the primary copy of the data or based on a combination thereof.

As another example, applying the updates to the redundant data 212 at the secondary facility 104 may be performed in response to a trigger received from the application at the computer 108 that originated the updates. Alternately, a system administrator may initiate the trigger. In still another example, updates may be based on an external clock-driven event. For example, updates may occur periodically, once each day, week, month, or year. Updates may occur upon certain specified times and dates. Further, a combination of techniques may be applied. For example, a batch of operations may be applied to the redundant data after a specified time interval unless a possible irregularity in the data has been detected through a data consistency check.

If a possible irregularity has been detected, further updates to the redundant data may be halted until further investigation is performed, such as by a system administrator. Accordingly, multiple batches may be queued at the secondary facility 104. In the event that the write queue 122 fills up, further updates to the primary copy at the primary facility 102 may be blocked. Alternately, rather than blocking the write transactions, the transactions may be stored at the primary facility 104 (e.g., as a single large group); if even that is insufficient, the transactions may simply be remembered in a manner that requires a fixed, known amount of space (e.g., by a bitmap-like structure of updated blocks, tracks, segments, or cylinders), and updates to the primary copy allowed to proceed. For example, a system administrator may select between blocking the updates and storing them at the primary facility.

In one aspect, the size of the receive batches m , $m+1$, $m+2$, etc. may be determined according to time intervals. For example, new receive batches may be started at specified time intervals. These time intervals may be the same as or different from any time interval used for delaying application of a batch. Alternately, the size of the receive batches may be determined according to the predetermined quantity (e.g., by a number of transactions or send batches or by storage capacity consumed) to be included in the receive batch. By increasing the size of the receive batches and/or the amount of time they are queued at the secondary facility, this will tend to increase the opportunity for preventing errors from propagating to the redundant data. However, this will also tend to increase the size of the queue needed in the secondary facility 104 which will tend to increase its cost. Accordingly, a trade-off can be made based on cost and the potential adverse consequences of error

propagation. Further, the receive batch size may be adjusted adaptively, such as based on the available space for the write queue 122 in the secondary facility 104. Thus, to conserve space by allowing more overwrites, the batch sizes may be increased.

5 As described, a single write queue 116 and 122 may be present at each of the primary facility 102 and the secondary facility 104. In which case, write transactions directed to different data objects, such as files or logical units (LUs), may be queued together. Alternately, multiple write queues may be maintained at either or both of the primary and secondary facilities 102 and 104. For example, a separate write
10 queue may be associated with each file being updated or with each LU, or with a “consistency group” of LUs that must be updated consistently).

Figure 6 illustrates an exemplary flow diagram of a method 300 for queuing and applying a batch of transactions at a secondary storage facility 104 in accordance with an aspect of the invention. Performance of the steps of the method 300 may be
15 performed under control of the secondary controller 118 (Figure 3). In step 302, one or more write transactions are received into the write queue 122 (Figure 3) at the secondary facility 104. As mentioned, the write transactions may be received one at a time or in groups (e.g., n, n+1, n+2, etc.). In step 304, the operations are preferably collected into batches (e.g., m, m+1, m+2). This may include replacing an earlier
20 operation with a later-received operation that affects the same data. As shown in Figure 2, this step includes sending write records and corresponding data to the log 210.

In step 306, a determination is made as to whether the current batch is complete. As mentioned, this determination may be based, for example, on a time
25 interval for collecting operations into the batch or upon the number of operations or quantity of data to be included in the batch. If the batch is not complete, program flow may return to step 302 for collecting additional operations as needed to complete the batch. Once the batch is complete, program flow moves from the step 306 to a step 308. Meanwhile, a subsequent batch may be formed in the same manner.

30 In step 308, a determination may be made as to whether the completed batch is ready to be applied to the redundant data at the mass-storage media 120. As mentioned, this determination may be based on elapsed time, a specified event (e.g., a data consistency check) or a combination thereof. If the batch is not ready to be applied, program flow may remain in the step 308 until the batch is ready to be

applied. Note that if an excessive time elapses, a timeout error may be indicated in step 308 or if a check of the data to be applied indicates an irregularity, a data integrity error may be indicated in step 308. When an error is indicated, the process applying batches at the secondary facility 104 is preferably halted until the source of the error is resolved. As mentioned, under these circumstances, transactions may be halted at the primary facility 102 or may be stored at the primary facility 102.

Assuming it is determined in step 308 that a batch is ready to be applied (i.e. committed) to the redundant data 212 (Figure 2), the batch is applied in step 310. Meanwhile, the determination of step 308 may be made relative to a subsequent batch. In this manner, multiple batches are successively queued in the secondary storage facility 104 and applied to the redundant data at the secondary storage facility. As shown in Figure 2, data for a batch is applied by sending it to the LU 212. As also shown in Figure 2, once the data for a batch (e.g., a send batch) has been applied, the secondary 104 may send an acknowledgement to the primary 102.

Thus, an asynchronous redundancy technique has been described in which write transactions are queued at a secondary storage facility so as to inhibit propagation of errors, for example, in the event of a software error at a primary storage facility, and so as to minimize loss of data in the event of a failure at the primary storage facility.

Figure 7 illustrates a state machine 400 for controlling the operation of the primary data storage facility 102 and/or the secondary data storage facility 104, in accordance with an aspect of the invention. The state machine 400 of Figure 7 may be implemented, for example, by the appliances 202 and 204 of Figure 2 which may include appropriately configured hardware, software or firmware in disk arrays, storage devices, hosts (e.g., computer 108), in-host I/O bus adapters, network switches, network hubs, or combination thereof, which may be dedicated or may have shared functionality.

In a preferred embodiment, the state machine 400 controls operation of a local data storage facility, while a duplicate instance of the state machine 400 controls operation of a remote storage facility. Because both facilities may be controlled by state machines having substantially the same set of states, only one state machine 400 is illustrated in Figure 7. It will be apparent, however, that two or more such state machines 400, provided at local and remote sites, may be operative at any one time.

The state machine 400 is divided generally into two regions, as shown by the horizontal dotted line in Figure 7, depending upon whether the facility is acting as a primary facility (e.g., 102 of Figure 1) or as a secondary facility (e.g., 104 of Figure 2). More particularly, the states above the dotted line control operation as a primary
 5 facility, while the states below the dotted line control operation as a secondary facility.

Assuming the facility is acting as a primary facility, and under normal operating conditions (e.g., in absence of a fault at the primary facility), operation is controlled by a “normal” state 402 (such state names are merely exemplary). If the
 10 facility is acting as a secondary facility under normal operating conditions, operation is controlled by a “normal” state 404. When the local and remote facilities are both in their normal states 402 and 404, respectively, the system 100 may operate generally as described above in which updates are forwarded from the primary facility 102 to the secondary facility 104.

15 Certain faults may occur with respect to a primary facility 102. These include, for example, the primary log 206 becoming filled to a predetermined capacity, a failure of the storage device(s) that hold the primary log 206, a failure of the storage device(s) that hold the primary copy 208 of the data, a failure which renders the local facility inoperative, such as a failure of the appliance 202, or a
 20 failure that renders the remote facility inaccessible to the local facility or inoperative, such as a failure of the storage device(s) that hold the secondary log 210 or the redundant data 212, a communication failure (e.g., in medium 106 of Figure 1) or a failure of the appliance 204.

After such a fault, one or more recovery events may occur. For example, after
 25 a failure of the primary log 206, the primary log 206 may become operational again, such as by repair or replacement of a failed storage device that stores the log 206. Also, after a fault at the remote facility or a fault that renders the remote facility inaccessible to the local facility, the remote facility may be returned to service. Upon returning to service, the remote facility may still contain its redundant copy of the
 30 data 212 and the secondary log 210 or the remote facility may be treated as empty of data.

Other fault and recovery events may occur with respect to the secondary facility 104. Possible faults include, for example, the secondary log 210 becoming filled to capacity, or a failure that causes the local facility (acting as the primary 102)

to cease sending updates to the secondary copy 212, or a failure of the remote facility, such as a failure of the storage device(s) that hold the redundant data 212 or a failure of the appliance 204. Possible recovery events include, for example, returning the remote facility to service. Upon returning to service, the remote facility may still
 5 contain its redundant copy of the data 212 and the secondary log 210 or the remote facility may be treated as empty of data.

Referring again to Figure 7, when the local facility (which was operating in normal state 402) experiences a fault so that it is essentially inoperative, it ceases acting as the primary 102. This is illustrated in Figure 7 by a “failed” state 406
 10 (which may be entered via transition 408). In addition, the remote facility may cease acting as the secondary 104 and, instead, the remote facility enters a “failover” state 410 from its normal state 404 (via transition 412). The secondary facility 104 may not detect when the primary 102 has failed since this may appear the same to the secondary 104 as though the primary 102 is simply quiet. Thus, entry into the
 15 failover state 410 may require intervention, for example, by a system administrator after the fault at the primary 102 has been discovered. Alternately, certain failures of the primary facility 102 may be detected, for example, by the primary 102 and the secondary 104 periodically exchanging status or keep-alive messages. If the primary facility 102 fails to send one or more expected messages or sends a message
 20 indicating failure has occurred, the secondary 104 may recognize that a fault has occurred at the primary 102 so that it may automatically take action.

In the failover state 410, the remote facility prepares to function as the primary facility 102. This includes the remote facility committing any data in its secondary log 210 to the redundant data 212. During the failover state 410, write requests from
 25 the computer 108 may be paused during which time the computer 108 queues the requests. The remote facility then assumes the role of the primary 102 so that request traffic from the computer 108 is redirected to the remote facility. Redirecting the traffic may be accomplished, for example, by the remote facility sending an appropriate notification to the computer 108; alternately, one or more other host
 30 computers may assume the role of computer 108 after the failover.

If the local facility has not recovered by the time the remote facility assumes the role of primary 102, the remote facility enters a standalone state 414 from the failover state 410 (via transition 416). In the standalone state 414, the primary facility 102 appends new entries to its primary log 206, and accesses and updates data 208.

However, because the local (now: secondary) facility has been determined to be unavailable, the new entries are not propagated to the secondary 104.

Thus, a technique has been described in which state machines are employed to cause a remote facility to assume the role of primary in the event of a fault affecting the local facility. Changes in roles between the facilities can be in response to other
 5 the local facility. Events (referred to herein as “failover” events), such as a fault affecting the remote facility or a fault affecting a communication medium between the facilities, or an operational condition, such as a manually initiated event (e.g., a system administrator initiating the change in roles) or an automatically initiated event (e.g., the change is
 10 prearranged to occur at a particular time), or in response to communication traffic conditions (e.g., a greater portion of request traffic originating closer to the second data storage facility – explained in more detail herein).

Eventually, the local facility may recover. Assuming the local facility becomes functional again, it preferably resumes operation as the secondary 104.
 15 However, before resuming operation as the secondary 104, the local facility preferably attempts to ensure that its data is consistent with that in the remote facility (acting as the primary 102). More particularly, the local facility determines whether it still has its copy of the data intact (now, the redundant data 212) and, if so, whether its data is up-to-date with respect to the primary copy 208. This resumption of a
 20 previous role by one of the facilities may be referred to as a “fallback” event and may be performed in response to conditions other than a fault or a fault recovery (at the primary, the secondary or a communication medium between the primary and secondary), including those events described previously as failover events.

For example, recovery of the local facility may be detected by the remote
 25 facility (acting as the primary 102) if the local facility resumes sending keep-alive or status messages. In response, the remote facility (primary 102) may signal the local facility that the primary has updates in its primary log 206. Alternately, upon becoming functional, the local facility may send a request for updates to the remote facility to determine whether the primary log 206 at the remote facility (acting as the
 30 primary 102) is empty.

If the log 206 is empty, this indicates that the data at the local and remote facilities is consistent. If the local facility recovers with its data intact and there is no inconsistency, it may transition from the failed state 406 directly to the normal state 404 (via transition 418). In the normal state 404, the local facility functions as the

secondary facility 104. In addition, the remote (now: primary) facility may enter the normal state 402 from the standalone state 414 (via transition 420). Alternately, depending upon which state the remote facility was in, it may enter the normal state 402 from the failover state 410 (via transition 422). In normal state 402, the remote
5 facility functions as the primary facility.

However, if there are records in the primary log 206, this means there is an inconsistency between the data held at the local and remote facilities. Accordingly, the local facility may transition to a pending state 424 (via transition 426). In the pending state 424, a backup for the primary log 206 is forwarded to the local facility.

10 In addition, the remote facility may transition to the normal state 402 (via transition 420 or 422). The updates are sent to the log 210 and then committed to the redundant data 212 at the local (now: secondary) facility. Once these records are committed, the local facility may transition from the pending state 424 to the normal state 404 (via transition 428).

15 If the local facility was failed for an extended period of time or has lost its data (e.g., repairs may have required replacement of its storage devices with empty ones), the amount of data required to update the local facility before it can begin normal operation as the secondary 104 may be expected to exceed the capacity of its secondary log 210. Thus, the entire contents of the data to be stored redundantly (a
20 “snapshot” of the primary copy 208) may be sent to the local facility. In this case, the local facility (acting as the secondary 104) moves to a direct update state 430 (via transition 432). In addition, it may signal the primary facility 102 to enter a data propagation state 434 (via transition 436 or 438). In the data propagation state 434, the entire contents of the data to be stored redundantly (a “snapshot” of the primary
25 copy 208) may be sent from the remote facility to the local facility. This may include condensing the data, such as by using known techniques for data compression.

Then, the remote facility (operating in the state 434) sends the condensed data to the local facility (operating in the update state 430) which commits the data to the redundant version 212, preferably bypassing the secondary log 210. Once the entire
30 snapshot is committed to the redundant data 212, the remote facility may enter the normal state 402 (via transition 440), while the local facility may enter the normal state 404 (via transition 442).

As described, from the failed state 406, the local facility may move to the normal state 404, to the pending state 424, or to the update state 430, depending on

the circumstances. Also, from the failover state 410, the remote facility may enter the normal state 402, a standalone state 414, or the data propagation state 434, depending on the circumstances. Moreover, the local facility may move from the pending state 424 to the failover state 410 (via transition 444) in the event that the remote facility
5 experiences a fault before the local facility enters the normal state 404.

Once the remote facility has entered the normal state 402 and the local facility has entered the normal state 404, the facilities have exchanged roles. Thus, a technique has been described in which state machines are employed to exchange the roles of primary and secondary between local and remote facilities.

10 To change back, the two facilities commit all of the outstanding updates and then resume their original roles. This may be accomplished by the local facility, which was operating as the secondary 104 in normal state 404, transitioning to normal state 402 (via transition 446) and resuming functioning as the primary 102. Also, the remote facility, which was operating as the primary 102 in normal state 402,
15 transitions to the normal state 404 (via transition 478) and resumes functioning as the secondary 104. Request traffic from the computer 108 is also redirected to the local facility.

Returning the local facility to its role as primary 102 and returning the remote facility to its role of secondary 104 is preferably performed when traffic between the
20 two facilities is quiet. This may be accomplished by an application that is running on the computer 108 initiating the role reversal during a period that the computer 108 does not require access to the primary facility 102. Alternately, either of the facilities 102, 104, may signal the other and the computer 108 to initiate the role reversal.

Thus, a technique has been described in which state machines are employed to
25 return the local and remote facilities to their original roles.

As described, the exchanging of roles may be performed in response to a fault, or it may be initiated in response to other conditions. For example, the origin of storage request traffic within the system 100 may be used to reverse the roles of the storage facilities. More particularly, in a distributed system, multiple host computers
30 108 at different locations may access the primary facility 102 for performing storage operations. During certain periods, the greatest portion of requests to the primary 102 may be originated by computers 108 that are physically closer to the secondary 104. Under these circumstances, efficiency would tend to be increased if the role of the primary facility was shifted closer to the origin of the communications as this would

shorten the communication distances (e.g., by reducing the amount of system traffic and communication latencies). Accordingly, the origins of storage requests may be monitored by the primary facility 102 (e.g., as an operation performed in the normal state 402). Based on relative locations of the local facility, the remote facility and the
 5 origins of the requests, the primary facility 102 may determine that it would be more efficient to shift its role to the remote facility or to the local facility. In response, the roles may be shifted, as needed, depending upon the current traffic patterns.

During normal operation in which the primary facility 102 is in state 402 and the secondary facility 104 is in state 404, a fault may occur in which the secondary
 10 facility 104 becomes inoperative. Under these circumstances, the secondary facility 104 may enter the failed state 406. Upon recovery, the secondary facility 104 may return to the normal state 404 directly, or via the pending state 412 or the update state 414, as described above.

From the perspective of the primary facility 102, a fault at the secondary
 15 facility 104 or a communication failure between the primary 102 and secondary 104 may result in the secondary 104 becoming unavailable to the primary 102. These faults may be detected by periodically exchanging keep-alive or status messages between the secondary 104 and the primary 102 during normal operation. Absence of the messages from the secondary 104 indicates the secondary 104 is unavailable. In
 20 response, the primary facility 102 may enter the standalone state 414 from the normal state 402 (via transition 448).

If the secondary 104 recovers while the primary 102 is in the standalone state 414, the primary 102 may return to the normal state 402 (via transition 420). The updates logged at the primary facility 102 may then be forwarded to the secondary
 25 facility 104 (while the secondary 104 is in pending state 424). However, if the primary log 206 becomes filled to capacity before the secondary facility 104 becomes available, the primary facility 102 may transition from the standalone state 414 to a bitmap state 450 (via transition 452).

In the bitmap state 450, the primary facility 104 effectively condenses the
 30 logged records using a bitmap or other type of change record. Thus, the bitmap takes the place of the entries in the log 206. The bitmap is a record that includes an indication for each data block of the primary copy 208 that has changed. For example, the bitmap may include a logical “zero” or “one” for each data block of the primary copy, where a logical “one” indicates that the block has changed since it was

last propagated to the secondary. Thus, unlike the log 206, which includes the changes to the data, the bitmap only indicates whether the data for a block has changed. The write-ordering of the changed portions is generally not preserved by the bitmap. While the primary 102 is in the bitmap state 450, any new updates are

5 incorporated into the bitmap. Because write-ordering is not preserved, the bitmap specifies a large batch of updates that preferably are be committed to the redundant data 212 at the secondary 104 as a whole. Thus, if the secondary 104 recovers while the primary is in the bitmap state 450, the primary 102 transitions to a data

10 propagation state 434 (via transition 456). In this state 434, the data blocks indicated by the bitmap are propagated to the secondary 104, where they are preferably committed to the redundant data 212 as a whole. If the secondary 104 becomes unavailable again while primary 102 is in the data propagation state 434, propagation of the data is halted and the primary 102 returns to the bitmap state 450 (via transition 454) where the primary 102 continues to incorporate new updates into the bitmap.

15 When the primary 102 is in the standalone state 414 because the secondary 104 is unavailable, the secondary 104 may recover without its redundant data. For example, either or both of the secondary log data 210 or the redundant data 212 may be lost if the secondary facility 104 is replaced or repairs required replacement of its storage devices. Similarly to the situation described above in which the local facility

20 resumes operation as secondary 104 after a fault, under these circumstances, the entire contents of the primary copy 208 may need to be sent to the secondary 104. Thus, primary facility 102 transitions from the standalone state 414 to the data propagate state 434 (via transition 438) in which a complete snapshot of the primary copy 208 is propagated to the secondary 104, if needed. Updating the secondary 104 occurs with

25 the secondary in the update state 430. Preferably, all changes are committed to the secondary 104 as soon as practical to limit any inconsistency with the data at the primary 102.

Thus, a technique has been described in which state machines are used to provide redundant data to a remote facility that experienced a fault.

30 While in the normal state 402, the primary log 206 may become filled though the secondary 104 is still accessible. For example, a burst of requests from the computer 108 or heavy traffic on the communication medium 106 between the primary 102 and secondary 104 may result in the primary log 206 becoming filled. In response, the primary facility 102 may transition to the data propagation state 434 (via

transition 458) in which the primary log 206 may be condensed to a bitmap and the corresponding data updates propagated to the secondary 104. As a result, the primary log 206 is emptied. Meanwhile, the secondary 104 may transition from its normal state 404 to the update state 430 (via transition 460). When propagation of the data is complete, the primary facility 102 may return to the normal state 402 from the data propagation state 434 (via transition 440) and the secondary facility 104 may return to the normal state 404 (via transition 442).

In the event that the primary log 206 fails, the primary facility 102 may enter a logless state 462. The logless state 462 avoids use of the primary log 206 and may be entered from any of the other states for the primary facility 102 (via transition 464). The logless state 462 is useful when a storage device which holds the primary log 206 fails or otherwise becomes inaccessible. In the logless state 462, the primary copy 208 is updated in response to write requests from the computer 108; however, the redundant data 212 is not updated since the primary log 206 is not available for this purpose.

To recover after the primary log 206 becomes available, the entire snapshot of the primary copy 208 may be propagated to the secondary 104 by the primary 102 transitioning to the data propagation state 434 (via transition 466). The secondary 104 transitions to the update state 430 (via transition 460). Then, operation may resume in the normal states 402 and 404. However, if the primary log 206 recovers before any updates cause the redundant data 212 to become inconsistent, the propagation state 434 may be bypassed (via transition 468) so that operation resumes in normal state 402 or in the bitmap state 450 if the primary log 206 is filled upon its recovery (via transition 470). Further, if the secondary 104 should become unavailable by the time the primary log 206 recovers, then the primary 102 may transition to the standalone state 414 from the logless state 462 (via transition 472).

When the secondary 104 is in the normal state 404, its secondary log 210 may become filled. Under these circumstances, the secondary 104 transitions to the update state 430 in which updates are propagated directly to the redundant data 212, bypassing the secondary log 210. By so doing, the redundant data 212 may become inconsistent with that of the primary 102; however, this is generally preferable to preventing updates to the redundant data 212. A snapshot update may eventually be performed (e.g., in the state 430) to bring the secondary 104 back into consistency.

Under certain circumstances, faults may be essentially unrecoverable. This is shown by a disaster state 474 in Figure 7, which may be entered from any state by transition 476. For example, assume one of the facilities has failed and the remaining, non-failed facility in is the standalone state 414. If the remaining facility should fail
5 before the failed facility recovers, the system 100 may be left in a state from which it cannot recover without intervention (e.g., a system administrator may be able to repair the system) or in a state in which it is unable to respond to new write requests from the computer 108. Also, if the primary 102 facility fails while the secondary facility 104 is in the direct update state 430, intervention will likely be required to restore the
10 system 100 to operation.

Thus, state machines have been described for local and remote facilities, in which each facility is able to cope with a variety of events.

As explained above, when the primary 102 experiences a fault, applications can fail over onto the secondary facility 104, which becomes the new primary facility.
15 Data left in the primary log 206 prior to the fault occurring at the local facility is essentially considered lost and new data is written to the remote facility acting as the new primary facility. When the local facility recovers, the data in its LU and log may be inconsistent with that in the new primary facility. In other words, each of the facilities may have a piece of data that the other does not have.

20 In some circumstances, in response to this inconsistency, it may be desirable for a system administrator or application-level utility to attempt to reconstruct the data so as to minimize or eliminate data loss. In other circumstances, it may be desirable to accept the potential loss of data by undoing some of the updates. This eliminates the inconsistency and allows operation of the system to continue. In one aspect of the
25 invention, a technique is provided for automatically responding to such a data inconsistency, as explained below.

Figure 8 illustrates an example of update and back-up copy propagation during failover and recovery in accordance with an aspect of the present invention. As shown in Figure 8, data blocks are written to sequentially in time at a local storage
30 facility acting as the primary storage facility 102. This is shown by a sequence 502 including versions A0, B0, C0, A1, B1, C1 of data blocks A, B and C. Because the data block version A0 is the left-most block in the diagram, this means that the data block A was written to first. In the example, the data block B is written to next, as shown by version B0, and, then, the data block C, as shown by version C0. Next, the

data block A is written to again, as shown by version A1. The data blocks B and C are then written to again, as shown by versions B1 and C1, respectively. It will be apparent that the ordering of writing the blocks is exemplary and will depend on the applications that generate the storage activity.

5 Meanwhile, updates to the data blocks A, B, C are sent to the remote facility acting as the secondary facility 104. This is shown in Figure 8 by the sequence 504 including versions A0, B0, C0 of data blocks A, B and C. Some delay is associated with sending the updated versions A0, B0 and C0 to the remote facility; accordingly, they are shown as being stored at the remote facility later in time.

10 Then, assume that at time T_0 , the local facility experiences a failover event (e.g., a fault). In response, the remote facility assumes the role of primary 102. Thus, after time T_0 , new updates are only stored at the remote facility. This is shown by an update 506, including an updated version B2 of the data block B.

 Assume that the local facility recovers at a time T_1 . At the time that the fault
15 occurred (T_0), some of the versions A0, B0, C0, A1, B1 and C1 may have been committed to the primary copy 208 (Figure 2), though others may not. In addition, because of the fault, it may not be possible to determine which versions have been committed. Thus, after recovery (after the time T_1) the local facility sends a request to the remote facility for its current copy of committed data (e.g. committed data
20 blocks), referred to as a “backup” of the data. In the example of Figure 8, the backup batch 508 is shown to include versions A0, C0 and B2. The backup 508 includes A0 and C0 from the sequence 504 and B2 from the update 506. If no blocks had been written during the interval between T_0 and T_1 , the update 506 would be empty and the backup 508 would include versions A0, B0 and C0.

25 The backup 508 is preferably sent to the local facility as a single batch, which is committed at the local facility as a whole. As explained above in reference to Figure 4, propagation of the backup 508 may be performed while the local facility is in the pending state 424. Note that versions A1, B1 and C1 are omitted from the backup 508 since they are not present at the remote facility. Thus, at the time T_2 , the
30 data at both facilities is consistent, while some data has been lost (i.e. A1, B1 and C1). In particular, at time T_2 each facility has committed data 510, including versions A0, C0 and B2. Because the operations that result in data consistency may be performed under control of the state machines 400 (Figure 7), they may be performed automatically. For example, forwarding the update 508 may be performed in pending

state 424 of Figure 7, while the local facility is preparing to resume operation as the secondary. Thus, operation from a state of data consistency can proceed without requiring human intervention.

In another example, some additional data blocks may have been written after the local facility experienced the fault, but before its recovery. This is shown in Figure 9, in which versions D0 and E0 of data blocks D and E were also written between the time T_0 and T_1 . The backup is thus expanded to include all other blocks that were written to before the committed blocks A, B and C. This is shown in Figure 9 by the backup 508' which includes the versions A0, C0, D0 and B2. Because version D0 was written before B2 in update 506' of Figure 9, the backup 508' needs to include D0 and also needs to be committed as a whole to maintain data consistency. Note that version E0 can be omitted from this batch and sent later. This is because E0 was written after the version B2 in the update 508' and, thus, it can be propagated after the update 508', in a second update 508". Alternately, E0 could be included in the update 508'.

As shown in Figure 9, at time T_2 , the data at both facilities is consistent, while some data has been lost (i.e. A1, B1 and C1). In particular, at time T_2 each facility has committed data 510', including versions A0, C0, D0, B2 and E0. These operations that result in data consistency may be performed under control of the state machines 400 of Figure 7 and, thus, may be performed automatically without requiring human intervention.

Figure 10 illustrates a third example of update and back-up copy propagation as may occur during failover and recovery in accordance with an aspect of the present invention. As before, Figure 10 illustrates that updates to the data blocks A, B, C are sent to the remote facility acting as the secondary facility 104. This is shown in Figure 10 by the sequence 502 including versions A0, B0, C0, A1, B1, C1. Meanwhile, the update sequence 504 including versions A0, B0, C0 is sent to the remote facility acting as the secondary facility 104. Then, assume that at time T_0 , the local facility experiences a fault. In response, the remote facility assumes the role of primary 102. Thus, after time T_0 , new updates are only stored at the remote facility. This is shown by the update 506' in which versions D0, B2 and E0 are written between the time T_0 and T_1 .

Then, assume that at time T_1 , the local facility recovers. However, before the remote facility can forward the updates to the local facility, assume that the remote

facility experiences a fault. As shown in Figure 10, this occurs at time T_2 . In this case, the local facility will attempt to commit its logged updates instead of receiving an update from the remote facility. Assuming that the version A0, B0 and C0 of the sequence 502 were committed, but versions A1, B1 and C1 were not, versions A1, B1 and C1 may still be in the local log 206 (Figure 2). Thus, the local facility attempts to commit the versions A1, B1 and C1 of the data blocks A, B and C to the primary copy 208. Accordingly, at time T_3 , the primary facility may have committed data 514 including versions A1, B1 and C1. In contrast to the circumstances of Figure 9, the versions D0, B2 and E0 are lost in this case, whereas, the versions A1, B1 and C1 (which were lost in Figure 9) are recovered.

As before, these operations that result in data consistency may be performed under control of the state machines 400 of Figure 7 and, thus, may be performed automatically without requiring human intervention. For example, committing the data 514 may occur in the failover state 410 (Figure 7) as the local facility prepares to resume the role of primary.

Thus, failover techniques for ensuring consistency between the primary and secondary copies of data have been described for a data redundancy system.

As described herein, when a fault occurs at the local facility acting as the primary storage facility 102, the remote facility may assume the role of primary 102 and cease acting as the secondary storage facility 104. However, if a fault then occurs at the remote facility, this may cause the system 100 to be unable to respond to any new write requests from the computer 108. Thus, under some circumstances, it may be desired to provide additional fault tolerance.

In one embodiment, redundant components are provided at the local facility and/or the remote facility to provide additional fault tolerance. For example, a back-up appliance that is redundant of the primary appliance 202 (Figure 2) may be provided at the local facility. This will tend to prevent a fault at the appliance 202 from forcing the remote facility to take over for the primary. Similarly, an LU that is redundant of the primary LU 208 may be provided at the local facility. This will tend to prevent a fault at the data storage devices which hold the primary LU 208 from forcing the remote facility to assume the role of the primary.

Figure 11 illustrates primary and secondary storage facilities in which redundant elements are provided in accordance with an embodiment of the present invention. As shown in Figure 11, the primary facility 102 includes the appliance 202

for storing data in LU 208. In addition, a redundant appliance 202' is provided at the primary facility 102. Preferably, the redundant appliance 202' is disjoint from its counterpart 202, meaning that it preferably does not share common hardware or software elements, or common failure modes.

5 In normal operation, in absence of a failure at the appliance 202 or at the redundant appliance 202', the appliance 202 is active and acts as master, while the appliance 202' acts as a shadow or slave. Thus, information about changes made to the active redundancy appliance's LU 208 in response to a write operation requested by the computer 108 are propagated to the shadow appliance 202'. These updates are
10 preferably made synchronously. Accordingly, after the change has been forwarded to the shadow 202' and an acknowledgement received, the write operation may proceed (e.g., by writing to the log 206, to the mass storage 208 and returning an acknowledgement to the host). Write records may be forwarded to the shadow appliance in batches. Further, the acknowledgements may be sent by the shadow
15 appliance in batches.

 If the redundant appliance 202' is updated synchronously, it will have the same data and state information as the primary appliance 202 when a fault occurs. Accordingly, the shadowing redundancy appliance 202' may take over operation from the master as soon as it is needed (using the data in LU 208 and the shared log 206).

20 Though redundant logs may be maintained at each redundancy appliance the primary 102, both of the primary redundancy appliances 202 and 202' preferably have access to a common log 206. This is shown in Figure 11. The common log 206 may be accessible from both appliances 202 and 202', for example, via a storage area network (SAN). Maintaining a single log 206 has an advantage of avoiding any
25 latency associated with making two synchronous updates to a shadow log. Rather, maintaining the shadow appliance's state only incurs extra latency of transferring update messages to the shadowing redundancy appliance (e.g., across an interconnecting LAN between the appliances 202 and 202'). The log 206 can be made fault-tolerant by using a local redundancy scheme, such as RAID5 or mirroring.

30 Fail-over occurs when the appliance 202 experiences a fault and the redundant appliance 202' then assumes the role of master. When a fault occurs at the appliance 202, this may be detected by the redundant appliance 202'. For example, a fault may be detected by the appliance 202' when the appliance 202 ceases communication with appliance 202' or if the appliance 202 sends a fault notification to the appliance 202'.

In response, the redundancy appliance 202' no longer acts as a shadow, but instead responds to read and write requests directly to the computer 108.

Similarly to the primary 102 facility, the secondary facility 104 may also include the appliance 204 for storing data in the LU 212 that is redundant of the data at the primary facility 102. In addition, a redundant appliance 204' may be provided at the secondary facility 104 that is redundant of the appliance 204. Similarly to the primary facility 102, the redundant appliance 204' is preferably disjoint from its counterpart 204.

In normal operation, in absence of a failure at the appliance 204 or at the redundant appliance 204', the appliance 204 is active and acts as master, while the appliance 204' acts as a shadow or slave. Changes made to the active redundancy appliance 204 (e.g., in response to a send batch received from the primary 102), are propagated to the shadow appliance 204'. These changes are preferably made synchronously; accordingly, once the change has been forwarded to the shadow 204', an acknowledgement may be returned to the primary 102. Because the redundant appliance 204' is preferably updated synchronously, it will have the same data and state information as the primary appliance 204 when a fault occurs and may take over operation from the master as soon as it is needed (using the shared log 210).

Assuming all of the appliances 202, 202', 204 and 204' are present and operational, a write request (e.g., issued by the computer 108) may be received at the appliance 202. In response, the appliance 202 causes a write record to be written synchronously in the primary log 206. In addition, the corresponding data for the request is written to the primary copy of the data at LU 208. The appliance 202 forwards the write record to the shadow appliance 202'. In response, the shadow appliance 202' updates any local data it maintains, and sends an acknowledgement to the appliance 202. An acknowledgement may then be sent to the computer 108 indicating the request was successfully stored by the primary facility 102.

Independently, the primary LU 208 and the secondary LU 212 may each be independently locally replicated, and made to have one or more local, redundant, shadow copies. This additional redundancy is in addition to any internal redundancy provided by the storage devices on which the LUs are stored. In particular, the redundant shadow LU copies preferentially reside on different storage devices, so that the redundant copy is stored disjointly from its counterpart, meaning that they preferably do not share common hardware or software elements or common failure

modes. Updates to the main copy LU 208 or LU 212 are preferably propagated to their shadow copy in the same order that they are written to the main copy; this may be achieved synchronously or asynchronously.

As described above in reference to Figures 3-5, the primary facility 102 may
5 batch the write requests into a send batch prior to forwarding them to the secondary facility 104. The appliance 204 at the secondary 104 then causes the batch to be written in the secondary log 210. In addition, the corresponding data for the batch is written to the primary copy of the data at LU 212, though committing the data to the LU 212 may be performed according to receive batches formed by the appliance 204
10 where the receive batches honor any receive barriers. The appliance 204 forwards the data to the shadow appliance 204'. This is preferably done by forwarding receive batches to the shadow appliance 204'. In response, the shadow appliance 204' records the writes of the data for the batch in LU 212 and sends an acknowledgement to the appliance 204. An acknowledgement may then be sent to the primary facility
15 102 indicating the request was successfully stored by the secondary facility 102.

While the primary 102 and secondary 104 are described as each having a redundant appliance and LU, this is not necessary. For example, the redundant appliance 202' and LU 208' may be provided at the local facility, but not at the remote facility. Thus, while the local facility acts as the primary, it will generally be
20 able to withstand more faults than otherwise before the remote facility needs to assume the role of primary.

As described, redundancy at the local and/or remote facilities is implemented by one of the appliances being active (i.e. acting as the master) and the other acting as the shadow (or slave). In an alternate embodiment, both appliances may be active. In
25 this case, each appliance 202 and 202' at the primary 102 may communicate directly with the computer 108 to perform operations (e.g., writes) in parallel. Similarly, each appliance 204 and 204' at the secondary 104 may communicate directly with the primary 102.

Thus, a data redundancy technique has been described in which redundant
30 components are provided at primary and/or secondary facilities for improved fault tolerance.

In one embodiment, multiple components are provided at the local facility and/or the remote facility for performing data storage operations in a distributed manner. Because the operations are distributed in such a facility, it is expected that

the facility will be able to accommodate a larger workload than otherwise (e.g., having a higher storage request rate or requiring additional storage capacity). For example, a pair of appliances may be provided at the local facility, each having a corresponding log and LU. Figure 12 illustrates primary and secondary storage

5 facilities in which data storage is distributed in accordance with an embodiment of the present invention. Thus, the appliance 202 uses the log 206 and the LU 208, while an additional appliance 202" uses an additional log 206" and LU 208". In this case, the logs and LUs store different data, such that the request workload to be serviced by the primary facility 102 is distributed among the appliances 202 and 202" and their

10 associated logs and LUs. Thus, some requests received by the primary facility 102 (e.g., from the computer 108) are serviced by the appliance 202, while others are serviced by the appliance 202". The requests may be assigned to the appliance 202 or the appliance 202" by the application that generates the requests (e.g., applications running on the computer 108) or by any other technique for assigning data to storage

15 units. For example, it may be desired to balance the loads between the two appliances 202 and 202" such as by assigning applications or data blocks to each such that each has an approximately equal workload.

So that the secondary facility 102 is able to accommodate the same workload as the primary facility 102, the secondary facility 104 is preferably also provided with

20 additional components corresponding to those of the primary facility 102. As shown in Figure 12, the secondary facility 104 includes the appliance 204 that uses log 210 and LU 212 and an additional appliance 204" that uses log 210" and LU 212".

In operation, when a write request is received at the appliance 202, a write record is written (preferably, synchronously) to the log 206 at the primary facility

25 102. In addition, the corresponding data for the request is written to a primary copy of the data in LU 208. An acknowledgement may then be sent to the computer 108 indicating the request was successfully stored by the primary facility 102. Similarly, when a write request is received at the appliance 202", a write record is written to the log 206" and the corresponding data is written to a primary copy of the data in LU

30 208". Once the data is successfully stored, an acknowledgement may then be sent to the computer 108.

Because the primary LUs 208 and 208" each store portions of the primary copy of the data at the primary facility 102, updates to the redundant version of the data at the secondary facility 104 are preferably coordinated among the appliances at

the primary 102 so as to maintain data consistency. For example, in absence of such coordination, in the event of a failure at one of the appliances 204 or 204", the other appliance may continue to function. Over time, the data in the LUs 212 and 212" may become inconsistent with the other since only portions of the data in the functioning

5 LU would continue to receive updates. This could result in inconsistency at the secondary facility 104, which could prevent application recovery in case of failure of the primary 102. Thus, in a preferred embodiment, updates are coordinated among the appliances at the primary 102 before the updates are forwarded to the secondary facility 104.

10 Coordinating updates among the appliances 202 and 202" at the primary 102 is preferably accomplished by coordinating send batch boundaries. In one aspect, one of the appliances 202 or 202" may first determine a time at which to terminate the current send batch according to any of the criteria described herein (e.g., by batch size or available communication bandwidth between the primary 102 and secondary 104).

15 The appliance at the primary 102 which determined the boundary may then declare a send batch boundary to the other appliance at the primary 102 by informing the other of the time at which its current batch should be terminated. In this way, each appliance 202 and 202" terminates its current batch at approximately the same point in time.

20 In another aspect, the batches may be terminated by the computer 108 declaring the end of a send batch. For example, an application at the computer 108 may issue a write operation that includes a predetermined end-of-batch indicator that is recognizable to the appliances 202 and 202". The computer 108 may, thus, signal both appliances 202 and 202" to end their current batches. In this manner, the

25 computer 108 may also specify in advance a time for the batches to end. As another example, the computer 108 may notify one of the appliances 202 or 202" to end the current batch. The appliance at the primary 102 that receives the notification may then notify the other appliance at the primary 102 to end its current batch. In this case, the computer 108 or the appliance that receives the notification from the

30 computer 108 may select a time at which the batches should be ended. Thus, terminating the send batches is a two-phase operation. In a first phase the end of batch is declared. In the second phase, each appliance 202 and 202" ends its current batch in accordance with the declaration.

Once the batches are terminated, each appliance 202 and 202" may then send its batch to a corresponding one of the appliances at the secondary facility 102. More particularly, data stored in LU 208 by appliance 202 at the primary 102 is forwarded to the appliance 204 for storage in the LU 212 at the secondary 104. In addition, data
5 stored in the LU 208" by appliance 202" at the primary 102 is forwarded to the appliance 204" for storage in the LU 212" at the secondary 104.

The appliances 204 and 204" at the secondary 104 then cause their respective batches to be written in the secondary logs 210 and 210" respectively. In addition, the corresponding data for the batches is written to the primary copy of the data at LU
10 212 and 212", respectively. Committing the data to the LU 212 may be performed according to receive batches formed by the appliances 204 and 204" where the receive batches honor any receive barriers.

In addition, to maintain data consistency at the LUs of the secondary facility 204, the appliances 204 and 204" preferably coordinate committing of the batches to
15 their respective LUs 212 and 212". Committing the batches at the secondary may also be performed as a two-phase operation. For example, one of the appliances 204 or 204" may first determine a time at which to commit the current batch according to any of the criteria described herein (e.g., after a data check). The appliance at the secondary 104 which determined the time to commit the batch may then declare this
20 to the other appliance at the secondary 102 by informing the other of the time at which its batch should be committed. In this way, each appliance 204 and 204" commits its batches to storage at approximately the same point in time. Well-known techniques such as "two-phase commit" can be used to ensure that both commit, or neither does. An acknowledgement may then be sent to the primary facility 102
25 indicating that the data was successfully stored by the secondary facility 104. The appliances 202 and 202" at the primary 102 may then delete the corresponding records in their logs 206 and 206". Deleting of the logs at the primary 102 can be performed independently (i.e. the appliances 202 and 202" need not coordinate deletion with each other).

30 Thus, a data redundancy technique has been described where functions at a primary and/or secondary facility are distributed for increased performance.

In embodiments described herein, the redundancy appliances 202 and 204 are each generally positioned in the path of write transactions to the LU associated with the appliance. For example, referring to Figure 2, the appliance 202 at the primary

facility 102 stores write transactions locally in the primary LU 208. In addition, the appliance 202 forwards the write transactions (e.g., in send batches) to the secondary facility 104. Thus, in some circumstances, the appliance 202 may become a bottleneck. To reduce this possibility, the functions of storing the write transactions locally and forwarding them to the secondary facility may be performed separately. For example, the primary appliance 202 may perform functions related to forwarding write transactions to the secondary facility 104, while the computer 108 (or other software and/or hardware at the primary facility 102) may perform operations related to local storage of data for write transactions.

While not necessary, this separation of functionality may be employed at both the local facility and at the remote facility. In this case, the local and remote facilities may each separate their functions of storing locally-originated write transactions locally and forwarding redundant data to another facility. Thus, when the remote facility assumes the role of primary (e.g., in the event of a fault at the primary or for transactions for which the remote facility in the primary), the remote facility may separate these functions in the same way as the local facility. Also, the remote facility may function as a primary storage facility for some data storage operations and as a secondary data storage facility for other data storage operations.

Figure 13 illustrates primary 102 and secondary storage facilities 104 in accordance with an embodiment of the present invention. As shown in Figure 13, the computer 108 (and possibly one or more additional computers 108) may store data in the log 206 and one or more LUs 208 without intervention by the appliance 202. Rather, the computer 108 (e.g. a host computer) includes a redundancy component 602 and a communication medium 604 for providing a communication path between the redundancy component and the primary log 206 and LU 208. As is also shown in Figure 13, the medium 604 allows communications between the computer 108 and the storage elements 206 and 208 to bypass the appliance 202. However, the appliance 202 is preferably coupled to the communication medium 604 for forwarding transactions to the secondary facility 104. The redundancy component 602 may be implemented as computer-based software, such as an I/O adapter or device driver and may include additional hardware at the computer 108. The communication medium 604 may be implemented as any kind of network, such as a storage area network (SAN) or a local area network (LAN).

Similarly to the primary facility 102, the secondary facility 104 may include communication medium 606 coupled to the secondary log 210 and LU 212. When the secondary 104 acts as a primary storage facility, the communication medium 606 allows communications between one more computer systems 108 to bypass the

5 redundancy appliance 204 at the secondary 104. The appliance 204 is preferably also coupled to the medium 606 for storing data received from the primary 102 in storage elements 210 and 212. The appliances 202 and 204 may communicate with each other by a network, for example, via a wide area network (WAN). The appliances 202 and 204 may be referred to as network portal redundancy appliances (or simply

10 “portals” or “appliances”) since they serve to connect the primary network 604 to the secondary network 606. It will be apparent that the redundancy portals 202 and 204 may be each be implemented as a dedicated device, or by a host computer (e.g. host 108) or by a storage device (e.g., a device that includes the LU 208).

In response to a write request issued by a computer 108, the redundancy

15 component 602 writes a record of the request to the log 206 including a copy of the data to be written. The redundancy component 602 also writes the data to the LU 208 and sends a reply to the application (e.g., at the computer 108) that issued the request. The portal 202 retrieves write records from the primary log 206 and forwards them to the secondary facility 104, while preserving the write-ordering of the requests.

20 So that the portal 202 is informed of the write-ordering of the requests in the log 206, each write record may have an associated timestamp. The portal 202 may then forward the write records to the secondary facility 104 in order of ascending timestamp values. In one embodiment, before storing a write record in the log 206, the redundancy component 602 obtains a timestamp for the request and stores the

25 timestamp in the log 206 as part of the record for the request.

In addition, so that the portal 202 is informed of the location of write requests in the log 206, the redundancy component 602 may notify the portal of the location of each write request after one or more write requests have been stored in the log 206. Retrieval and forwarding of the request may commence immediately after the data is

30 stored in the LU 208. Thus, the redundancy component 602 may also notify the portal 202 to retrieve the request as soon as it has been stored in the LU 208. In this case, the portal 202 may retrieve the request concurrently with the redundancy component 602 sending the reply to the requesting application.

The portal 202 may forward write requests according to send batches. In this case, the redundancy component 602 may overwrite records in the log 206 where a later received request overwrites an earlier request in the log 206. In this case, the redundancy component 602 informs the portal of the send batch boundaries so that the portal 202 may retrieve an entire send batch for forwarding to the secondary facility 104. In this case, the redundancy component 602 also preferably informs the portal of the location of receive barriers (see Figure 5 herein) so that receive batches may be committed at the secondary 104 as a unit. Where send batches are employed, the redundancy component 602 need only inform the portal 202 of parameters for batches, such as a timestamp and location of the batch, rather than informing the portal 202 of the timestamp and location for every write transaction in the log 206.

When the portal 204 at the secondary facility 104 receives a write record (or a batch of write records), the secondary portal 204 writes the record (or batch) into the secondary log 210 and may then send an acknowledgement to the primary portal 202. In response, the primary facility 102 may delete the corresponding records in the log 206. Preferably, the primary portal 202 deletes the records, though this may also be accomplished by the redundancy component 602.

The secondary facility 104 commits data (e.g., in receive batches) to the secondary LU 212 in the order received from the primary facility 102 and then deletes the corresponding records from the secondary log 210. Storing the data in the LU 212 and deleting the records from the log 210 are preferably performed by the portal 204 though a redundancy component 602 at the secondary 104 may perform some of these functions.

Thus, a data redundancy technique using host and portal components has been described.

It may be desired to provide a technique for testing operation of a data redundancy system. For example, in view of the number of states of the state machine 400 and the numerous possible event sequences, it would be desirable to verify operation of a data redundancy system using one or more the state machines 400. As explained above in reference, transitions among the states of the state machine 400 of Figure 7 and operation within the states occur in response to external events, such as write requests and faults (e.g., a log disk fault or a network fault) and in response to internal events, such as log space exhaustion, update propagation and batch commits. In one aspect, the system 100 employing one or more state machines

400 is tested by applying a sequence of events and evaluating behavior of the system 100 in response. This testing may be performed, for example, on a physical implementation of the system 100 or on a simulation of the system 100 (where the simulator uses software and/or hardware to simulate operation of the system 100).

5 Figure 14 illustrates a flow diagram of a method 700 for testing a data redundancy system 100 or its simulation (the two will be treated as equivalent in what follows) in accordance with an embodiment of the present invention. In a step 702, a sequence of one or more events is applied to the data redundancy system 100. The events may include, for example, write operations, failover or fallback events,
10 including failures of various elements of the system 100, such as a log disk being non-responsive or a network communication interruption, and various recovery events, such as a failed log disk becoming responsive again or a network becoming operational again after a communication interruption.

 In a step 704, one or more verifications are then performed on the system 100.
15 In a preferred embodiment, the verifications include one or more of the following: a state validity verification, a data consistency verification and a system liveness verification.

 The state validity verification determines whether, after receiving the event(s) of step 702, the current state of the system 100 is valid with respect to the state
20 machines 400. For example, the current state of the primary facility 102 and the secondary facility 104 may be checked to determine whether each is in a valid one of the states of its respective state machine 400. This may also include determining whether the two states are consistent with each other. For example, when the local facility is in the standalone state 414, the remote facility should generally be in the
25 failed state 406.

 The data consistency verification determines whether the primary and secondary copies of the data 208 and 212 are consistent (e.g., in the normal states 402 and 404) and that the sequence of updates at the secondary facility 104 is a prefix of the sequence at the primary 102 (i.e., the updates at the secondary 104 include updates
30 from the primary 102 in their proper order, although the secondary 104 may be missing updates that have not yet been propagated from the primary 102).

 The liveness verification determines, in absence of a failure, whether data for all write requests received at the primary 102 are eventually written to both the primary 102 and the secondary 104. To check liveness, external events (e.g., write

requests and failures) are ceased and the system 100 is allowed to continue operation until all events have been processed. Then, the system 100 is checked to ensure that all of the write requests generated (in step 702) have been processed at both the primary 102 and the secondary 104.

5 Once the verifications in step 704 have been performed, a determination is made in step 706 as to whether all the verifications were positive. If not, this indicates that an operational error was encountered. Accordingly, the system 100 may be evaluated (i.e. debugged by a human operator) in step 708 to identify and correct the cause of the error. Once this has been completed, a sequence of one or more
10 events may again be applied in step 702 and the verification(s) performed again in step 704. If all the verifications are satisfied in step 706, this indicates that system 100 operated correctly in response to the events applied in step 702. Thus, one or more new events may be applied to the system 100 in step 702 and the process of performing the verifications may be repeated.

15 Accordingly, testing of the system 100 is performed incrementally, in which the system 100 (or a simulation thereof) is injected with a sequence of external events until a verification fails. When a verification fails, the system 100 is debugged and the error corrected. This process may be repeated until a sufficient number of events and event sequences have been applied to ensure that the system 100 operates as
20 intended.

 Thus, a technique for verifying operation of a data redundancy system has been described.

 While the foregoing has been with reference to particular embodiments of the invention, it will be appreciated by those skilled in the art that changes in these
25 embodiments may be made without departing from the principles and spirit of the invention, the scope of which is defined by the appended claims.

What is claimed is:

- 1 1. A data redundancy system including a primary facility for
2 storing a primary copy of data and a secondary facility for storing data
3 that is redundant of the primary copy of the data, wherein the primary
4 facility comprises:
5 a first redundancy appliance for receiving a sequence of write
6 requests and for storing data for the sequence of write requests in
7 storage associated with the primary storage facility; and
8 a second redundancy appliance for shadowing the first
9 redundancy appliance wherein the second redundancy appliance
10 assumes the role of the first redundancy appliance.
- 1 2. The data redundancy system according to claim 1, wherein the
2 first redundancy appliance updates the second redundancy appliance
3 with information about the data in storage.
- 1 3. The data redundancy system according to claim 1, wherein the
2 second redundancy appliance assumes the role of the first redundancy
3 appliance in response to a failover event.
- 1 4. The data redundancy system according to claim 3, wherein the
2 failover event comprises a fault.
- 1 5. The data redundancy system according to claim 4, wherein the
2 fault occurs at a location selected from the group consisting of the first
3 storage facility, the second storage facility and a communication
4 medium between the first data storage facility and the second data
5 storage facility.
- 1 6. The data redundancy system according to claim 3, wherein the
2 failover event comprises an operational event.

1 7. The data redundancy system according to claim 6, wherein the
2 failover event is selected from the group consisting of a manually
3 initiated event and an automatically initiated event.

1 8. The data redundancy system according to claim 3, wherein the
2 failover event comprises a communication traffic condition.

1 9. The data redundancy system according to claim 8, wherein the
2 communication traffic condition comprises a greater portion of request
3 traffic originating closer to the second data storage facility.

1 10. The data redundancy system according to claim 1, wherein the
2 first redundancy appliance updates the second redundancy appliance
3 with information about the data in storage.

1 11. The data redundancy system according to claim 10, wherein
2 the first redundancy appliance updates the second redundancy
3 appliance by forwarding a write record to the second redundancy
4 appliance for each write request.

1 12. The data redundancy system according to claim 11, wherein
2 write records are forwarded to the second appliance in batches.

1 13. The data redundancy system according to claim 11, wherein
2 the second redundancy appliance sends an acknowledgement in
3 response to receiving the write record.

1 14. The data redundancy system according to claim 13, wherein
2 the first redundancy appliance sends an acknowledgement to an
3 originator of the write request in response to receiving the
4 acknowledgement from the second redundancy appliance.

1 15. The data redundancy system according to claim 11, wherein
2 acknowledgements are sent in batches.

1 16. The data redundancy system according to claim 1, wherein the
2 first redundancy appliance is disjoint from the second redundancy
3 appliance.

1 17. The data redundancy system according to claim 1, wherein the
2 second redundancy appliance maintains a shadow copy of data stored
3 by the first redundancy appliance wherein the second redundancy
4 appliance accesses the shadow copy of the data when the second
5 redundancy appliance assumes the role of the first redundancy
6 appliance.

1 18. The data redundancy system according to claim 1, wherein the
2 second redundancy appliance accesses the data stored in storage by the
3 first redundancy appliance when the second redundancy appliance
4 assumes the role of the first redundancy appliance.

1 19. The data redundancy system according to claim 1, wherein the
2 first redundancy appliance forwards redundant data to the secondary
3 facility according to send batches.

1 20. The data redundancy system according to claim 1, wherein a first write
2 transaction in a send batch at the first storage facility is replaced with a second
3 write transaction where the second write transaction overwrites data written by
4 the first write transaction.

1 21. The data redundancy system according to claim 20, wherein a send
2 batch size is determined according to a time interval for collecting transactions
3 into the send batch.

1 22. The data redundancy system according to claim 20, wherein a
2 send batch size is determined according to a predetermined quantity of
3 transactions to be included in the send batch.

1 23. The data redundancy system according to claim 20, wherein a
2 send batch size is determined adaptively based on available
3 communication bandwidth between the first storage facility and the
4 second storage facility.

1 24. The data redundancy system according to claim 1, wherein the
2 secondary facility comprises:
3 a third redundancy appliance for storing redundant data for the
4 first sequence of write requests in storage associated with the
5 secondary facility; and
6 a fourth redundancy appliance for shadowing the third
7 redundancy appliance wherein the fourth redundancy appliance
8 assumes the role of the third redundancy appliance in the event of a
9 fault at the first redundancy appliance.

1 25. The data redundancy system according to claim 24, wherein
2 the third redundancy appliance stores redundant data to the storage
3 associated with the secondary facility according to receive batches.

1 26. The data redundancy system according to claim 25, wherein
2 the third redundancy appliance stores a receive batch after the receive
3 batch is determined safe to apply.

1 27. The data redundancy system according to claim 26, wherein a
2 receive batch is determined safe to apply when a data integrity test is
3 successfully performed.

1 28. The data redundancy system according to claim 25, wherein a
2 first write transaction in a first receive batch at the second storage
3 facility is replaced with a second write transaction in a second receive
4 batch wherein the second write transaction overwrites data written by
5 the first write transaction and the first and second receive batches are
6 combined.

1 29. A data redundancy method for storing a primary copy of data
2 at a primary storage facility and for storing data that is redundant of
3 the primary copy at a secondary storage facility, the method
4 comprising:
5 receiving a sequence of write requests at a first redundancy
6 appliance at the primary storage facility;
7 storing data for the sequence of write requests in storage at the
8 primary storage facility;
9 updating a shadow redundancy appliance at the primary storage
10 facility with information about the data in storage; and
11 forwarding the sequence of write requests to the secondary
12 storage facility.

1 30. The data redundancy method according to claim 29, wherein
2 said first redundancy appliance performs said storing and arranging
3 prior to a failover event and wherein the second redundancy appliance
4 performs said storing and arranging after the failover event.

1 31. The data redundancy method according to claim 29, wherein
2 said updating is performed synchronously.

1 32. The data redundancy method according to claim 31, wherein
2 said updating comprises the first redundancy appliance forwarding a
3 write record to the second redundancy appliance for each write
4 request.

1 33. The data redundancy method according to claim 32, further
2 comprising the second redundancy appliance sending an
3 acknowledgement in response to receiving the write record.

1 34. The data redundancy method according to claim 33, further
2 comprising the first redundancy appliance sending an
3 acknowledgement to an originator of the write request in response to

4 receiving the acknowledgement from the second redundancy
5 appliance.

1 35. The data redundancy method according to claim 29, wherein
2 the first redundancy appliance is disjoint from the second redundancy
3 appliance.

1 36. The data redundancy method according to claim 29, further
2 comprising the second redundancy appliance maintaining a shadow
3 copy of data stored by the first redundancy appliance.

1 37. The data redundancy method according to claim 29, further
2 comprising the first redundancy appliance arranging the sequence of
3 write requests according to a sequence of send batches, wherein said
4 forwarding comprises forwarding the send batches to the secondary
5 facility.

1 38. The data redundancy method according to claim 37, further comprising
2 determining a send batch size according to a time interval for collecting
3 transactions into the send batch.

1 39. The data redundancy method according to claim 37, further
2 comprising determining a send batch size according to a
3 predetermined quantity of transactions to be included in the send
4 batch.

1 40. The data redundancy method according to claim 37, further
2 comprising determining a send batch size adaptively based on
3 available communication bandwidth between the first storage facility
4 and the second storage facility.

1 41. The data redundancy method according to claim 29, further comprising
2 replacing a first write transaction in a send batch at the first storage facility

3 with a second write transaction where the second write transaction overwrites
4 data written by the first write transaction.

1 42. The data redundancy method according to claim 29, further
2 comprising:

3 receiving write requests from the primary storage facility at a
4 third redundancy appliance at the secondary storage facility;
5 storing data for the write requests in storage at the secondary
6 storage facility.

7
8 43. The data redundancy method according to claim 42, further
9 comprising updating a shadow redundancy appliance at the secondary
10 storage facility with information about the data in storage at the
11 secondary storage facility.

1 44. The data redundancy method according to claim 29, further
2 comprising storing data for the write requests in storage at the
3 secondary storage facility according to receive batches.

1 45. The data redundancy method according to claim 44, further
2 comprising determining whether a receive batch is safe to apply before
3 storing the receive batch.

1 46. The data redundancy method according to claim 45, wherein
2 said determining whether a receive batch is safe to apply comprises
3 performing a data integrity test on data of the receive batch.

1 47. The data redundancy method according to claim 44, further
2 comprising replacing a first write transaction in a receive batch at the
3 second storage facility with a second write transaction wherein the
4 second write transaction overwrites data written by the first write
5 transaction.

Abstract

Techniques for performing data redundancy operations in a fault-tolerant manner. In one aspect, a primary data storage facility stores a primary copy of data and a secondary facility stores data that is redundant of the primary copy of the data.

- 5 The primary facility includes a first redundancy appliance that receives a sequence of write requests and stores data for the sequence of write requests in storage associated with the primary storage facility. A second redundancy appliance shadows the first redundancy appliance and assumes the role of the first redundancy appliance in the event of a fault at the first redundancy appliance. In this way, fault tolerance is
- 10 increased by the presence of the second, shadow appliance.

1/14




FIG. 1

Title: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

Inventor: Ji et al.

Atty. Dkt. No. 200309044-1

2/14




FIG. 2

Title: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE
Inventor: Ji et al.
Atty. Dkt. No. 200309044-1

3/14




FIG. 3

Title: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE
Inventor: Ji et al.
Atty. Dkt. No. 200309044-1

4/14




FIG. 4

5/14




FIG. 5A




FIG. 5B

6/14




FIG. 6

7/14




FIG. 7

8/14




FIG. 8

9/14




FIG. 9

Title: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE
 Inventor: Ji et al.
 Atty. Dkt. No. 200309044-1

10/14




FIG. 10

11/14




FIG. 11




FIG. 12

Title: FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE
 Inventor: Ji et al.
 Atty. Dkt. No. 200309044-1

12/14

10456033.060603

13/14




FIG. 13

14/14




FIG. 14

PATENT APPLICATION

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATIONATTORNEY DOCKET NO. 200309044-1

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

FAULT-TOLERANT DATA REDUNDANCY TECHNIQUE

the specification of which is attached hereto unless the following box is checked:

() was filed on _____ as US Application No. or PCT International Application
Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
			YES: _____ NO: _____
			YES: _____ NO: _____

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION NUMBER	FILING DATE

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879Place Customer
Number Bar Code
Label hereSend Correspondence to:
HEWLETT-PACKARD DEVELOPMENT COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Direct Telephone Calls To:

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Minwen JiCitizenship: CanadianResidence: Sunnyvale, CAPost Office Address: 862 Hydrangea Ct., Sunnyvale, CA 94086-8115Inventor's Signature Minwen JiDate 6/6/03

Rev 02/03 (DecPwr)

(Use Page Two For Additional Inventor(s) Signature(s))

Page 1 of 2

10456053.060603

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)**

ATTORNEY DOCKET NO. **200309044-1**

Full Name of # 2 Joint Inventor: Alistair Veitch Citizenship: New Zealand
Residence: Mountain View, CA
Post Office Address: 2420 Whitney Dr., Mountain View, CA 94043-4133
Inventor's Signature: *A. C. Veitch* Date: June 6, 2003

Full Name of # 3 Joint Inventor: John Wilkes Citizenship: United Kingdom
Residence: Palo Alto, CA
Post Office Address: 3788 Park Blvd., Palo Alto, CA 94306-3430
Inventor's Signature: _____ Date: _____

Full Name of # 4 Joint Inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 5 Joint Inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 6 Joint Inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 7 Joint Inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

Full Name of # 8 Joint Inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature: _____ Date: _____

PATENT APPLICATION SERIAL NO. _____

U.S. DEPARTMENT OF COMMERCE
PATENT AND TRADEMARK OFFICE
FEE RECORD SHEET

06/10/2003 WASFAM1 00000059 002025 10456053

01 FC:1001	750.00 CH
02 FC:1202	486.00 CH

PTO-1556
(5/87)

PATENT APPLICATION FEE DETERMINATION RECORD

Effective January 1, 2003

Application or Docket Number

200309044-1

CLAIMS AS FILED - PART I

	(Column 1)	(Column 2)
TOTAL CLAIMS	47	
FOR	NUMBER FILED	NUMBER EXTRA
TOTAL CHARGEABLE CLAIMS	47 minus 20 =	* 27
INDEPENDENT CLAIMS	2 minus 3 =	*
MULTIPLE DEPENDENT CLAIM PRESENT ☐		

* If the difference in column 1 is less than zero, enter "0" in column 2

CLAIMS AS AMENDED - PART II

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT A	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
	Total *	Minus **	=
	Independent *	Minus ***	=
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM ☐		

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT B	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
	Total *	Minus **	=
	Independent *	Minus ***	=
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM ☐		

	(Column 1)	(Column 2)	(Column 3)
AMENDMENT C	CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA
	Total *	Minus **	=
	Independent *	Minus ***	=
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM ☐		

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.

** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20."

*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3."

The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

SMALL ENTITY TYPE ☐

OR OTHER THAN SMALL ENTITY

RATE	FEE		RATE	FEE
BASIC FEE	375.00	OR	BASIC FEE	750.00
X\$ 9=		OR	X\$18=	486
X42=		OR	X84=	
+140=		OR	+280=	
TOTAL		OR	TOTAL	1236

SMALL ENTITY

OR OTHER THAN SMALL ENTITY

RATE	ADDITIONAL FEE		RATE	ADDITIONAL FEE
X\$ 9=		OR	X\$18=	
X42=		OR	X84=	
+140=		OR	+280=	
TOTAL ADDIT. FEE		OR	TOTAL ADDIT. FEE	

RATE	ADDITIONAL FEE		RATE	ADDITIONAL FEE
X\$ 9=		OR	X\$18=	
X42=		OR	X84=	
+140=		OR	+280=	
TOTAL ADDIT. FEE		OR	TOTAL ADDIT. FEE	

RATE	ADDITIONAL FEE		RATE	ADDITIONAL FEE
X\$ 9=		OR	X\$18=	
X42=		OR	X84=	
+140=		OR	+280=	
TOTAL ADDIT. FEE		OR	TOTAL ADDIT. FEE	