

Preface

When I presented my experiences with the Event Replicator for Adabas at the International Users' Group Conference in Prague, 2010, Software AG and conference attendees inspired and encouraged me to write a book about replication to share my knowledge.

This book will discuss the basics of the Event Replicator and show how to set-up files for replication and create the definitions for destination, subscription, GFB, initial state and the new ADADBS function.

Not all utility functions can currently be automatically replicated on a non-mainframe platform. An example is file changes that require special handling on the subscription, the replicator engine and the target site. My book will walk you through the process step by step.

From a disaster and recovery standpoint, copying databases to the same or other hardware platforms becomes increasingly important. This book explains the difference between mirroring and replicating and shows the pros and cons of both techniques.

In tests with production-sized data of 160 million records, communication problems occurred, parameters weren't calculated correctly and RPL buffers and SLOG overflowed. The initial-state procedure was too time-consuming and a workaround had to be found.

In production with 1,400 updates per second, the source database showed higher CPU time and batch jobs ran much longer when replication was turned on.

Despite users' statements that the replication works correctly and comparisons between source and target are not necessary, self-written comparison programs sometimes show differences. Three methods will be discussed how to compare source and target files to ensure that both are identical.

Since database data is available in local or wide area networks and on public Web sites, security issues become even more important. One chapter in this book explains the available security software protecting the components during the replication process.

Last but not least, this book will also show the very important recovery function. Normally, the Event Replicator Server restarts automatically after an abnormal end and is able to recover any lost replication. But in some cases, data is lost and the replication must be replayed.

My special gratitude goes out to Larry Frazin, who supported me with replication tests. John Donnelly, Albert Stetson, Victor Tisuela, Hazel Baeza, Gary Walker and Chris Jonas who sent parameter values from their organizations. Roseanna Torretto monitored resources on the mainframe with RMF and Strobe. The team from Software AG, Wolfgang Weiss, Becky Albin, Alex Burggraf and Wil Heynen, who were always available for tuning discussions and product support. My thanks also to Michelle McGowan for her proofreading skills and emotional support.

Dieter W. Storr
Sacramento, California, U.S.A, August 2011

Contents

Preface	V
Table of Figures	XIII
1 Replication Basics	1
1.1 Today's Challenge	1
1.2 Replicating or Mirroring	3
1.2.1 Mirroring	3
1.2.2 Replication	4
1.3 Reasons for Replication	5
1.4 Check Your Intentions Before Starting	7
1.4.1 Mirroring or Replication?	7
1.4.2 Usage of Replicated Database	8
1.4.3 Measure/Compute and Determine Your Needs	9
1.4.4 Avoid Mistakes	10
2 Different Ways to Replicate	13
2.1 General Requirements / Basics	13
2.2 From Adabas to Adabas on the Mainframe	14
2.3 From Adabas on the Mainframe to Adabas on Linux, UNIX or Windows (LUW)	15
2.4 From Adabas on the Mainframe to Non-Adabas Databases	17
3 Definitions and Parameters	19
3.1 Getting Started	19
3.2 Source or Subscription Database (V8.1.3)	22
3.2.1 Workload Manager (WLM)	23
3.2.2 ADARUN Parameters	23
3.2.3 ADADBS Replicate Files	24
3.3 Replicator Engine – Reptor (V8.1.3/V3.2.1)	25
3.3.1 Workload Manager (WLM)	25
3.3.2 SLOG File	26
3.3.3 Replication System File	26
3.3.4 ADARUN parameters	27
3.3.5 Event Replicator Initialization Parameters	29
3.3.6 Define the Global Format Buffer	31
3.3.7 Define the Destinations	35
3.3.8 Define the Subscriptions	36
3.3.9 Define the Transaction Filter Definitions	37
3.3.10 Define the Initial-State	38

3.4	Unload and load parameters	38
3.5	Entire Net-Work (WCP V6.2.1) – Mainframe	41
3.5.1	ADARUN parameters	41
3.5.2	Parameter Statements	42
3.6	Entire Net-Work (WCP V7.3.3) – LUW	43
3.6.1	Parameters	43
3.6.2	Service.config	43
3.6.3	KZ008.KERNEL	44
3.6.4	Xts.config	44
3.7	Entire Net-Work (WCL V1.3.1) Client	44
3.7.1	Windows Service	44
3.7.2	UNIX Daemon	45
3.7.3	Xts.config	45
3.7.4	Service.config	45
3.7.5	Directory Server – xtsurl.cfg	46
3.8	Adabas on Windows (V6.2.1.1)	46
3.8.1	General	46
3.8.2	Parameters	46
3.9	EntireX Broker (V8.0.1)	48
3.9.1	EntireX Broker (V7.3.4)	48
3.9.2	Event Replicator Target Adapter	52
3.9.3	Schema and Replicated Data	53
4	Compare Files of Source and Target	57
4.1	Compare Amount Fields	58
4.2	API USR4011N – User Exit	59
4.3	Compare Highest ISNs (TOPISN)	60
5	Recovery After a Crash	63
5.1	Replay with a Token	64
5.2	Replay Without a Token	67
5.3	Replay-only Mode	68
5.4	Initial-state	68
6	Problems and Solutions	71
6.1	Source or Subscription Database	71
6.1.1	Replicator Buffer (LRPL) Overflow	71
6.1.2	High CPU Usage of Source DB	72
6.1.3	ADAFRK – Transactions Turned Off	79
6.1.4	Source DB didn't come down	79
6.1.5	File Changes Won't Replicate	80
6.2	Replicator Engine (Reptor)	86
6.2.1	NAB Overflow	86
6.2.2	Replicator Buffer (LRPL) Overflow	87

6.2.3	SLOG Overflow	89
6.2.4	RSP 153 from the Target DB	90
6.2.5	RSP 113 from the Target DB	92
6.2.6	RSP 98 from the Target DB	94
6.2.7	RSP 162 from the Target DB	95
6.2.8	Special characters are not correctly replicated	95
6.2.9	Wrong Replicated File in SLOG	96
6.2.10	RSP 148 From Target DB	97
6.2.11	Node is unreachable	100
6.2.12	Communication Problem	100
6.2.13	SLOG data disappeared	101
6.2.14	Replication is too slow	101
6.2.15	Initial state is too slow	102
6.2.16	New ADAFRI Message	103
6.2.17	Reptor under Stress	104
6.2.18	Replication with wrong FDT/GFB	105
6.3	Target Database	105
6.3.1	RSP 52 from the Target DB	105
6.3.2	RSP 176 from the Target DB	106
6.3.3	RSP 98 followed by RSP 113	106
6.3.4	RSP 113 on several files	107
6.3.5	RSP 153 – Displayed on Reptor	108
6.3.6	RSP 162 – LBP Space Problems	109
6.4	Utilities	112
6.4.1	Replay with more than 60 files	112
6.4.2	Replay with ERROR-139	112
6.4.3	Replay with ERROR-140 RSP 254	113
6.4.4	Replay with ERROR-140 RSP 131 Subcode 70	113
6.4.5	Replay with ERROR-140 RSP 131 Subcode 61	114
6.4.6	Replay with ERROR-148 – Token	114
6.4.7	Replay Abended with S0C4	115
6.4.8	Replay with Response 131 Subcode 61	117
7	Monitors	119
7.1	What can be monitored?	119
7.1.1	Source DB on the Mainframe	121
7.1.2	Replicator on the Mainframe	121
7.1.3	Entire Net-Work (WCP) on the Mainframe	124
7.1.4	Network – Line or Circuit	126
7.1.5	Firewalls, Routers, Load Balancing Systems	127
7.1.6	Entire Net-Work (WCP) on Windows	128
7.1.7	Adabas on Windows – Command Log	129

7.2	Self-Written	129
7.2.1	RPL and NAB Buffer Overflow	129
7.2.2	SLOG check on Reptor	132
7.2.3	Compare two SLOGs and Open Destination	134
7.2.4	SLOG Check for Logged Files	138
7.2.5	SLOG Space and Number of Records Check	139
7.2.6	Check Messages on DDPRINT	141
7.2.7	Target Check – Ping	142
7.2.8	Destination Report	144
7.2.9	Subscription Report	146
7.2.10	RPL Flag Check in ADAREP	148
7.2.11	Analyze WCP Log	151
7.2.12	Determine Highest ISN Currently Used	152
7.2.13	Compare File Components After Replication	152
7.3	Optimize for Infrastructure (O4I)	153
7.3.1	General	153
7.3.2	Dimensions	155
7.3.3	KPIs of Event Map AdabasEventRep	155
7.3.4	Built-In Rules	156
7.3.5	Dimensions – Destination	156
7.3.6	KPIs of Event Map AdabasEventRepDest	156
7.3.7	Built-In Rules – Destination	157
7.3.8	Dimensions – Input Queue	157
7.3.9	KPIs of Event Map AdabasEventRepInQueue	157
7.3.10	Built-In Rules	158
7.3.11	Dimensions – Subscription	158
7.3.12	KPIs of Event Map AdabasEventRepSubscr	158
7.3.13	Built-In Rules	158
7.3.14	How to Monitor a Component	159
8	Security	163
8.1	Source Database Adabas	163
8.1.1	Adabas SAF Security (ADASAF)	163
8.1.2	Adabas Security (ADASCR)	164
8.1.3	Adabas Data Encryption (CIPHERING)	164
8.1.4	Adabas Multi-client Files	165
8.2	Replicator Adabas	165
8.2.1	Adabas SAF Security (ADASAF)	165
8.3	Entire Net-Work (WCP) Mainframe V6.2.1	165
8.3.1	Entire Net-Work SAF Security (NETSAF)	165
8.3.2	Encryption V1.1.2	166
8.4	EntireX Broker	166
8.4.1	EntireX Security	166

8.5	Network	166
8.6	Entire Net-Work (WCP) LUW V7.3.4	167
8.6.1	Encryption V1.1.2	167
8.7	Entire Net-Work (WCL) Client V1.3.3	167
8.8	Event Replicator Target Adapter	167
8.9	Adabas ADASCR on Linux, Unix, Windows	167
8.10	Natural Security	168
8.10.1	NSC	168
8.10.2	Natural SAF Security	168
8.11	Non-Adabas on Linux, Unix, Windows	168
9	Appendix	169
9.1	Subscription DB – ADARUN parameters	169
9.2	Reptor – ADARUN parameters	171
9.3	Definitions Target DB – Windows	173
9.3.1	Parameters – db251.ini	173
9.3.2	High Water Marks	175
9.3.3	Activities	175
9.3.4	Buffer Pool Statistics	176
9.4	Test Results	177
9.5	Natural Monitor Programs	180
9.5.1	Ping Target Database	180
9.5.2	Check Messages on DDPRINT	181
9.5.3	RPL and NAB Buffer Overflow	183
9.5.4	SLOG Check on Reptor	185
9.5.5	Compare two SLOGs	186
9.5.6	Analyze WCP Log	189
9.5.7	Checksum Program	190
9.5.8	Compare Amount Fields	192
9.5.9	Parse Adabas Report for RPL Flag	194
9.5.10	Determine Highest ISN Currently Used	197
9.5.11	RPL30REP – SLOG Check	200
9.5.12	READF30P – Read SLOG Content	203
9.5.13	RPL-P002 – Compare File Components	204
9.6	WCP Log Windows	206
9.7	ADARPL Replay S0C4 Dump	213
9.8	Assembler Programs	216
9.8.1	Sleep/Wait	216
9.8.2	Get Dataset Name	218
9.9	EntireX V8.0.1 Parameters	220
	Bibliography	223
	Index	225