Informix Database Migrations, Exports & Imports (Part 2)

Mike Walker
mike@advancedatatools.com

Webcast on June 5, 2019

Advanced DataTools

Mike Walker

Mike Walker has been using Informix databases for over 20 years, as a developer and as a database administrator. Mike is an IBM Information Champion.

Mike runs the Remote DBA Support for Advanced DataTools Corporation.

mike@advancedatatools.com
www.advancedatatools.com
703-256-0267

Webcast Guidelines

- The webcast is being recorded. The webcast replay and slides will be available in a few days.
- Please mute your Line - background sounds will distract everyone.
- Use the Chat button in the toolbar at the bottom of the screen to ask questions.

Recap of Part 1

- Migrate to a new server:
  - Create target server
  - Create database & tables
  - Copy data (unload/copy/load)
  - Create indexes/constraints/views/procedures/ permissions/etc
  - Statistics
  - Validate
- Database Logging
- dbexport/dbimport

Speeding things up

- Running in parallel
  - Running multiple unload/load operations at the same time
  - Unload/Load parts of the same table/file at the same time
- Different utilities for unload and load

Unload & zip/unzip & Load

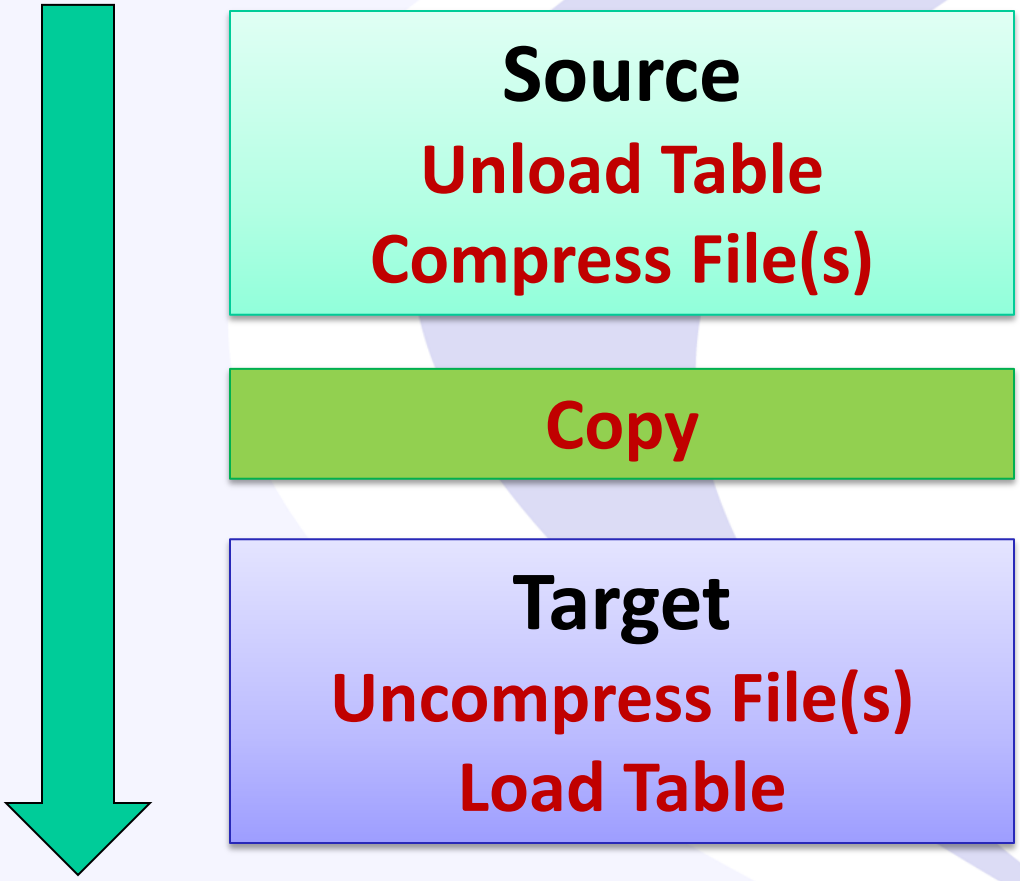

Unload & zip/unzip & Load

- Perform zip/unzip at the same time as unload/load
- Instead of files, use a pipe

```
mkfifo pipe.1
```

```
cat pipe.1 | gzip > file.unl.gz
```

- Unload to the pipe instead of file
- gzip will run against contents of pipe

```
-rw-r--r-- 1 informix informix   2295 Oct  3 09:44 file.unl.gz
```

Unload & zip/unzip & Load

```
mkfifo pipe.2
```

```
cat file.unl.gz | gunzip > pipe.2
```

- Load from the pipe instead of file
- gunzip will run against contents of pipe as the data is being loaded

Note that gzip/gunzip can use 100% of a single core, which may be a bottleneck. pigz is a parallel zip, but monitor the CPU used

Unload & Load Simultaneously

Do not wait for all the unloads to complete before loading

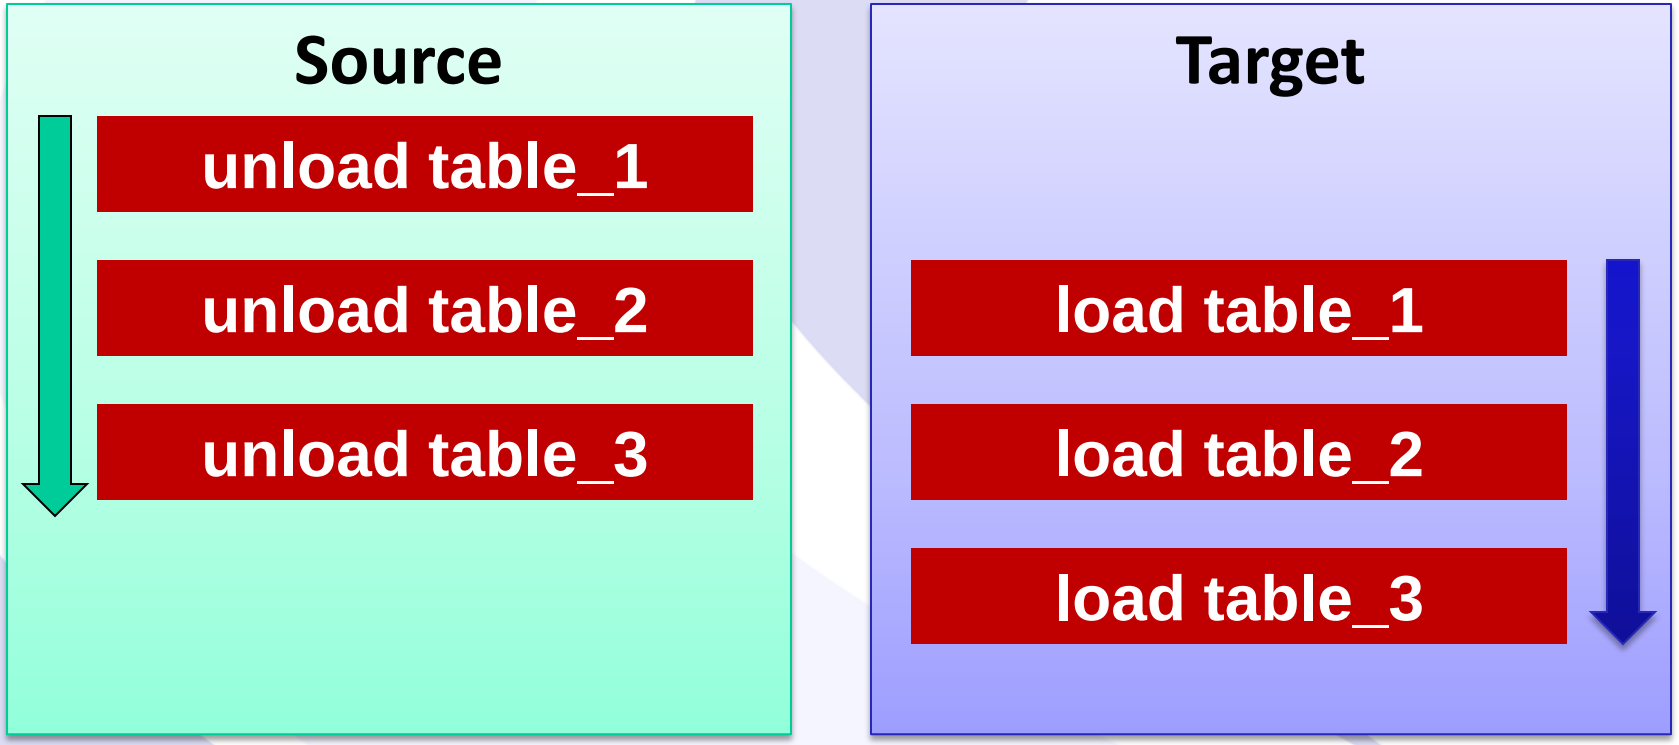

Unload & Load Simultaneously

Segment large tables, so loads can start while the remainder of the table is still being unloaded

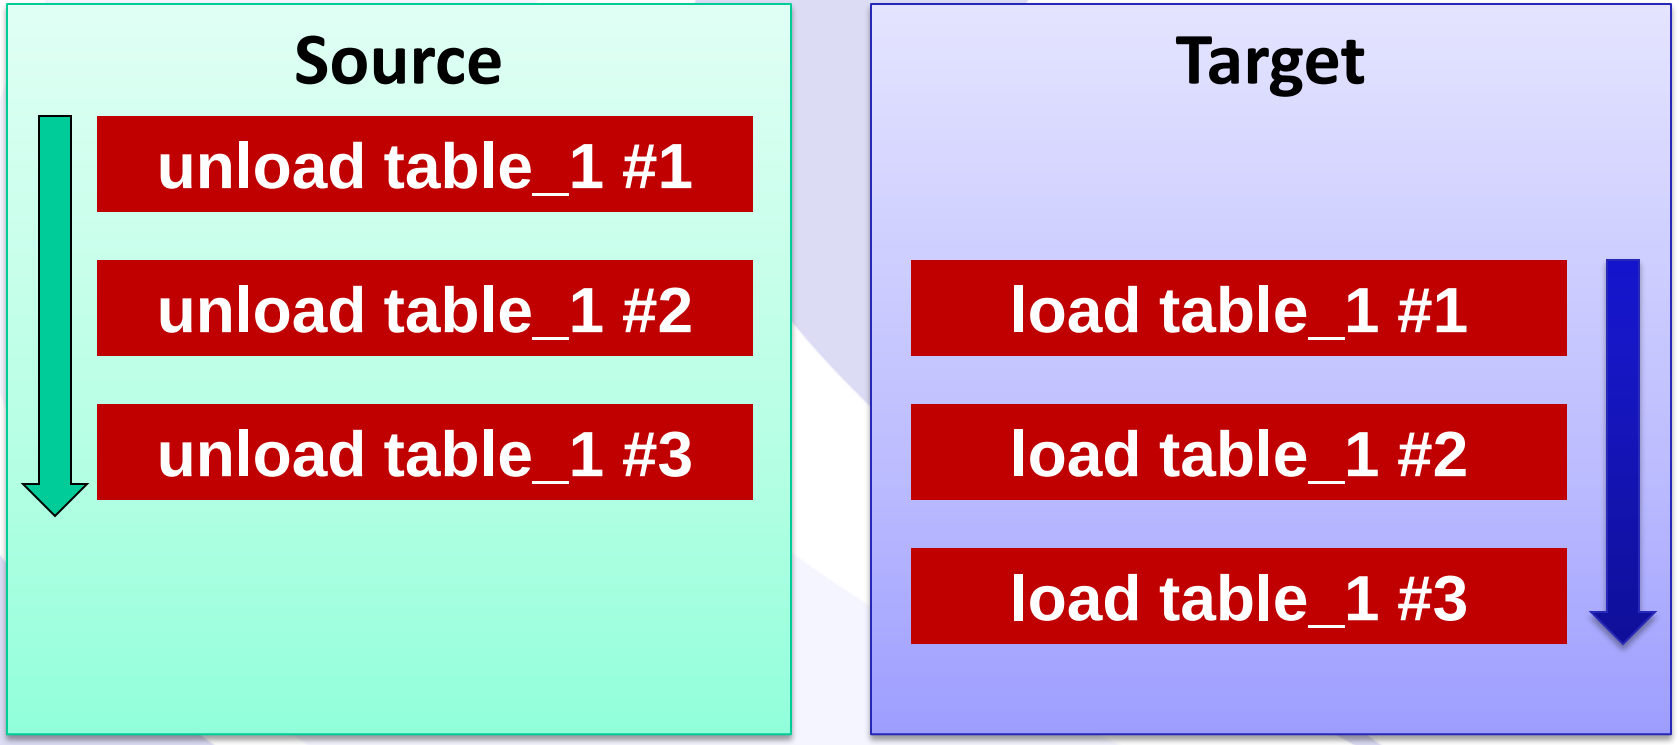

Unload & Load Multiple Tables

Unload/Load ***multiple*** tables ***at the same time***

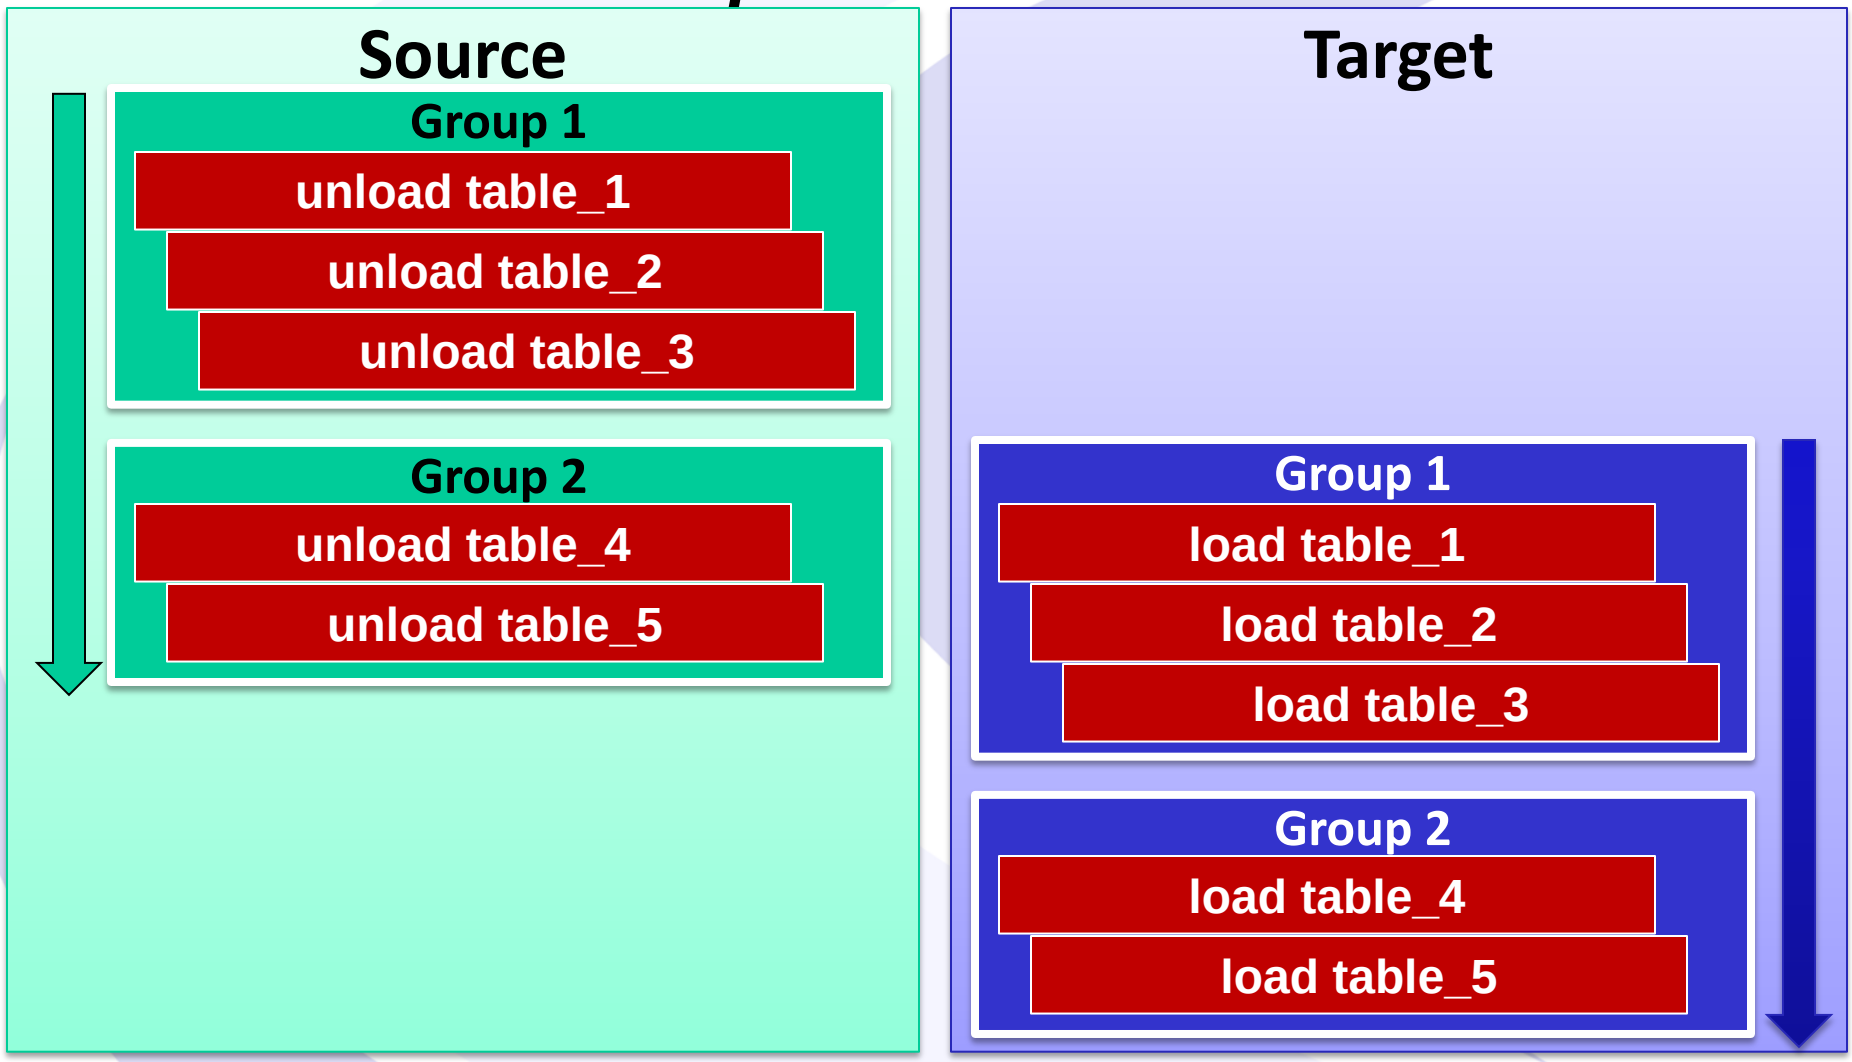

Unload & Load Multiple Tables

Unload/Load ***multiple partitions at the same time***

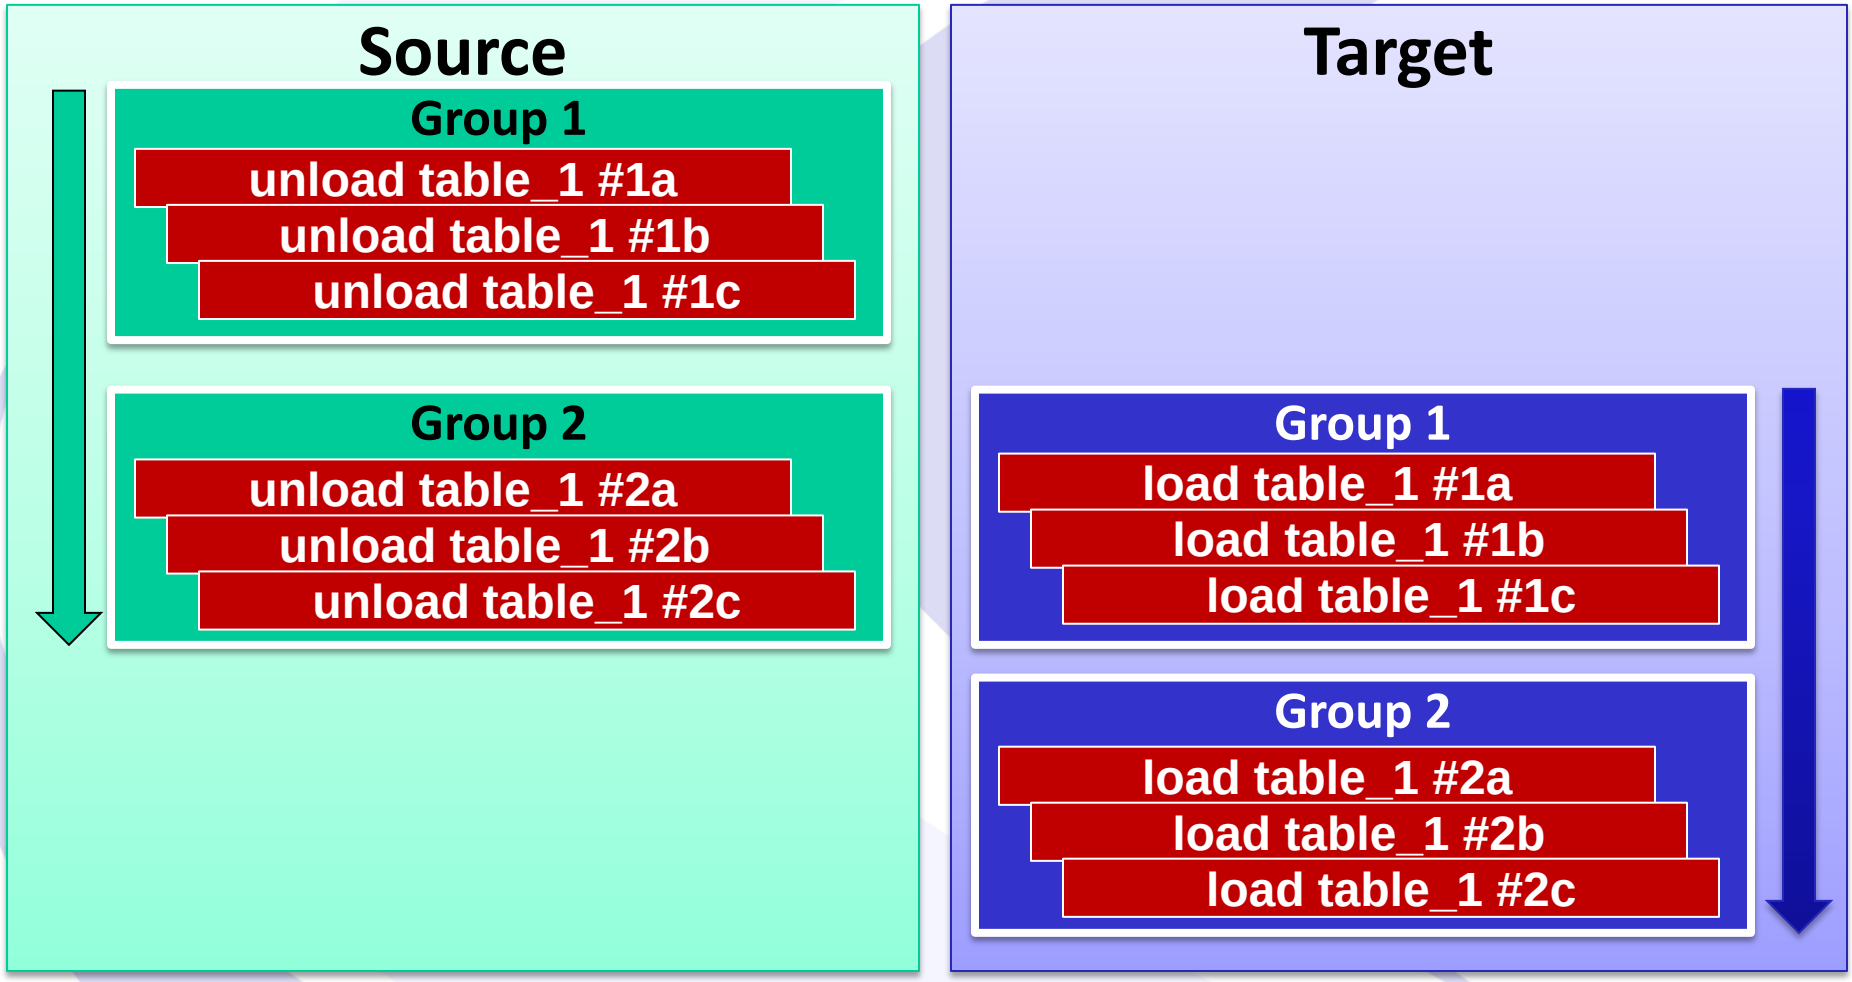

Plan & Test

- Running multiple steps in parallel is a logistical headache
  - Document the process
  - Script the process
- Organize tables in parallel group so all take a similar time to complete
- Organize groups to minimize the time that the target is idle
- Test to see how much parallelism the systems can handle, and adjust accordingly

Run multiple processes

- Script a process to run multiple steps, each as a background process
- Use the “wait” command so that script does not terminate until all steps are complete
- Each process should write to its own log file

Run multiple processes

```
function unload_tab {

   TAB=$1

   echo "Unloading Table: ${TAB} at " `date +"%D %T"`

   <do unload here>

   echo "Completed at" `date +"%D %T"`
}

echo "Starting unloads..."

(time unload_tab table1) > unl_table1.out 2>&1 &
(time unload_tab table2) > unl_table2.out 2>&1 &
(time unload_tab table3) > unl_table3.out 2>&1 &

echo "Waiting for unloads to complete..."
wait
echo "Unloads completed"
```

Don't forget the error checking!

What are the Largest Tables?

- Identify the largest tables
  - Determine the best utility for unload/load
  - Determine grouping/splitting
- Table ***size*** is more important than the ***number of rows***
- May decide to fragment tables on the target

What are the Largest Tables?

```
database sysmaster;

select
   n.dbsname dbname,
   n.tabname,
   sum(i.ti_nrows) nrows,
   sum(round(i.ti_npused *
          (i.ti_pagesize/1024))) used_kb,
   count(*) frags
from systabnames n, systabinfo i
where n.partnum = i.ti_partnum
  and n.tabname != "TBLSpace"
group by 1,2
order by 4 desc
```

What are the Largest Tables?

```
dbname    sales
tabname   order_line
nrows     349103815
used_kb   26971840
frags     1

dbname    sales
tabname   order_line_idx1
nrows     0
used_kb   21901840
frags     1

dbname    sales
tabname   order_header
nrows     348575727
used_kb   19944300
frags     1
```

Will include indexes also

Which Tables are Fragmented?

- Fragmented/Partitioned tables

  (not available in all Informix versions)

  – Greater benefits from using PDQ

  – Benefit from simultaneous unload/load processes

  – Benefit from unload/load utilities which can use multiple files

Which Tables are Fragmented?

```
select
   t.tabname,
   f.strategy,
   f.exprtext,
   f.dbspace,
   i.ti_nrows::int nrows,
   round(i.ti_npused *
          (i.ti_pagesize/1024))::int ptn_kb
from
   <database>:sysfragments f,
   <database>:systables t,
   sysmaster:systabinfo i
where f.tabid = t.tabid
  and f.fragtype = "T"
  and f.partn = i.ti_partnum
```

Which Tables are Fragmented?

```
tabname    orders
strategy   E
exprtext
((month = '09' ) OR (month = '10' ) )
dbspace    datadbs16
nrows      59238
ptn_kb     47408

tabname    orders
strategy   E
exprtext
((month = '11' ) OR (month = '12' ) )
dbspace    datadbs16
nrows      52247
ptn_kb     41824

tabname    orders
strategy   E
exprtext
remainder
dbspace    datadbs16
nrows      268649
ptn_kb     214992
```

insert into…select from

- Copy from remote system, insert into local
- Add remote server information to sqlhosts
- Reference remote table as:

  <database>@<server>:<table>

```
insert into orders
select *
from stores_demo@orangetcp:orders
```

- Logging modes on local and remote databases must match
- Simple, but performance is poor

SQL Unload/Load

Unload

```
unload to "<filename>"
select *
from <table>;
```

- Simple, but not particularly fast
- Only one output file
- Can unload to a named pipe
- ASCII/pipe-delimited (can change delimiter)
- May see slight improvement with setting PDQPRIORITY and isolation of DIRTY READ

SQL Unload/Load

Unload

```
set isolation to dirty read;
set pdqpriority 2;

unload to /migrate/orders.unl
select *
from orders;
```

orders.unl

```
1001|05/20/2008|104|express|n|B77836|06/01/2008|20.4|10.0|07/22/2008|
1002|05/21/2008|101|PO on box; deliver to back door
only|n|9270|05/26/2008|50.6|15.3|06/03/2008|
1003|05/22/2008|104|express|n|B77890|05/23/2008|35.6|10.8|06/14/2008|
1004|05/22/2008|106|ring bell twice|y|8006|05/30/2008|95.8|19.2||
1005|05/24/2008|116|call before
delivery|n|2865|06/09/2008|80.8|16.2|06/21/2008|
```

SQL Unload/Load

Load

```
load from "<filename>"
insert into <table>;
```

- Simple, but not very fast
- Only one input file
- Can load from a named pipe
- No automatic commits. Beware of:
  - Locks
  - Long transactions
- Alter table to raw, or use non-logging database
- Disable indexes & constraints while loading

SQL Unload/Load

Load – logged table/database

```
set pdqpriority 50;

set constraints for orders disabled;
set indexes for orders disabled;

begin work;

lock table orders in exclusive mode;

load from /migrate/orders.unl
insert into orders;

commit work;

set indexes for orders enabled;
set constraints for orders enabled;
```

May encounter a long transaction

SQL Unload/Load

Load – RAW table

```
set pdqpriority 50;

set constraints for orders disabled;
set indexes for orders disabled;

alter table orders drop constraint r120_39;
alter table orders drop constraint u120_36;

alter table orders type(raw);

load from /migrate/orders.unl
insert into orders;

set indexes for orders enabled;
set constraints for orders enabled;

alter table orders type(standard);

alter table orders add constraint primary key (order_num)
constraint orders_pk;

alter table orders add constraint foreign key (customer_num)
references customer constraint orders_cust_fk;
```

RAW tables can have no PK/FK constraints/unique idx, so must drop and recreate

For migration, unlikely to have indexes & RI constraints created

Name constraints

SQL Unload/Load

Load – Unlogged Database

```
set pdqpriority 50;

set constraints for orders disabled;
set indexes for orders disabled;

load from /migrate/orders.unl
insert into orders;

set indexes for orders enabled;
set constraints for orders enabled;
```

No need to lock table

No need to drop PK/FK/unique idx

dbload

- Command line utility
- Load only – not fast
- Uses a command file, containing:
  - Name of data file to load and format
  - Table to load
- Specify commit intervals – avoid long tx!
- Flexible (fixed format/delimited, limit bad rows, ignore first *n* rows, etc)
- Disable indexes & constraints, use non-logged table, etc before running dbload

dbload

Command File:

bigtab_load.cmd (delimiter format)

```
FILE /migrate/bigtab.unl DELIMITER '|' 2;
INSERT INTO bigtab_load;
```

bigtab_load_fixed.cmd (fixed format)

```
FILE /migrate/bigtab.unl
   (col1 1-10,
    col2 11-20);
INSERT INTO bigtab_load(id, data) values (col1, col2);
```

dbload

Command:

```
dbload -d stores -c bigtab_load.cmd -n 50000 -r -e 0
-l bigtab_load.err
```

- **-n :** Commit interval
- **-r :** Do not lock the table – allows multiple dbloads to be run against the same table
- **-l :** Log bad records
- **-e :** Number of bad records permitted

See documentation for all options

```
DBLOAD Load Utility      INFORMIX-SQL Version 12.10.UC9DE

50000 Row(s) loaded so far to table bigtab_load.

100000 Row(s) loaded so far to table bigtab_load.
.
.
950000 Row(s) loaded so far to table bigtab_load.

1000000 Row(s) loaded so far to table bigtab_load.

Table bigtab_load had 1000000 row(s) loaded into it.
```

External Tables

Introduced in IDS 11.50.xC6

- Define an external table based on files (or pipes)
- Insert into an external table with SQL writes to the files(s)/pipe(s) *[use for unloading to a file]*
- Select from an external table with SQL reads from the file(s)/pipe(s) *[use for loading from a file]*
- Can use multiple files/pipes – multi threaded
- Flexible (delimited/fixed format, max errors, etc)
- Fastest way to ***unload*** data (*possibly…*)
- Informix (internal) format is very fast
- Increase the FIFO VPs to improve performance
- See earlier webcast at: http://advancedatatools.com/Informix/Webcasts.html

External Tables

Unload

```
set pdqpriority 25;

create external table bigtab_ext
sameas bigtab
using (
   datafiles (
      "DISK:/migrate/bigtab.unl1",
      "DISK:/migrate/bigtab.unl2"));

insert into bigtab_ext
select * from bigtab;
```

External Tables

Unload – internal format

```
set pdqpriority 25;

create external table bigtab_ext
sameas bigtab
using (
   datafiles (
      "DISK:/migrate/bigtab.unl1",
      "DISK:/migrate/bigtab.unl2"),
   format "informix");

insert into bigtab_ext
select * from bigtab;
```

The internal format *may* not be compatible with the target system

External Tables

Load – internal format

```
set pdqpriority 25;

drop table if exists bigtab_ext;

create external table bigtab_ext
sameas bigtab_load
using (
   datafiles (
      "DISK:/migrate/bigtab.unl1",
      "DISK:/migrate/bigtab.unl2"),
   format "informix");

alter table bigtab_load type (raw);

insert into bigtab_load
select * from bigtab_ext;

alter table bigtab_load type (standard);
```

Assumes no PK/FK

Disable indexes/constraints if they exist

External Tables

Load – Deluxe vs Express

"Mode" only applies to ***loading*** through external tables

Express (faster):

- Can only be used on non logged tables
- Target table must have no indexes

Deluxe (slower):

- Evaluates indexes and unique constraints during load
- Data can be accessed during the load
- Cannot specify a commit interval – long tx possible & many locks!

External Tables

Load – internal format/express

```
create external table bigtab_ext
sameas bigtab_load
using (
   datafiles (
      "DISK:/migrate/bigtab.unl1",
      "DISK:/migrate/bigtab.unl2"),
   format "informix",
   express);

alter table bigtab_load type (raw);

insert into bigtab_load
select * from bigtab_ext;

alter table bigtab_load type (standard);
```

```
onbar -b -F
```

Backup required before can modify the table

High Performance Loader

- HPL utility can be used for unloads and loads
- Very flexible, but can be finicky
- Supports multiple unload/load files or pipes
- Conversion/No-Conversion (ascii/internal format)
- Deluxe/Express mode
- Fastest way to ***load*** data (*possibly*…)
- No-conversion files may be compatible with external tables “informix” format

High Performance Loader

Create unload job – no-conversion

```
onpladm create job order_dtl_u -d /tmp/order_dtl.unl -D
sales -t order_dtl -fu -n

Successfully created Job order_dtl_u
```

High Performance Loader

Run unload job

```
onpladm run job order_dtl_u -fu
Connecting to onpload, Please wait...
Successful connection to onpload established
Fri May 31 18:25:57 2019

SHMBASE         0x0000000044000000
CLIENTNUM       0x0000000049010000
Session ID 5

Unload Database -> sales
Query Name      -> AUTO.5
Device Array    -> order_dtl_u
Query Mapping   -> AUTO.5
Query           -> select * from 'informix'.order_dtl  for read only
Convert Reject  -> /tmp/order_dtl_u.rej
18:25:59 Records Processed ->  10416
18:25:59 Records Processed ->  20832
  <snip>
18:25:59 Records Processed ->  145824

Database Unload Completed -- Unloaded 155471 Records  Detected 0 Errors
Fri May 31 18:25:59 2019

Job Completed Successfully ... connection closed
```

High Performance Loader

Create load job – no-conversion, express

```
onpladm create job order_dtl_l -d /tmp/order_dtl.unl
-D sales -t order_dtl -fl -n

Successfully created Job order_dtl_l
```

High Performance Loader

Run load job

```
onpload -j order_dtl_1 -i 50000

19:36:32  Requested shared memory segment size rounded
from 4096KB to 4580KB
```

```
onbar -b -F
```

Backup required before can modify the table

High Performance Loader

/tmp/order_dtl_l.log

```
cat /tmp/order_dtl_1.log
Fri May 31 19:36:36 2019

SHMBASE         0x0000000044000000
CLIENTNUM       0x0000000049010000
Session ID 8

Load Database   -> sales
Load Table      -> order_dtl
Device Array    -> order_dtl_l
Record Mapping  -> AUTO.8
Convert Reject  -> /tmp/order_dtl_1.rej
Filter Reject   -> /tmp/order_dtl_1.flt
19:36:39 Records Processed ->  50024
19:36:39 Records Processed ->  100048
19:36:39 Records Processed ->  150072
Table 'order_dtl' will be read-only until level 0 archive

Database Load Completed -- Processed 155471 Records
Records Inserted-> 155471
Detected Errors--> 0
Engine Rejected--> 0

Fri May 31 19:36:40 2019
```

BLOBs

- **Distributed Query:**
  - Cannot insert Smart BLOBs (BLOB/CLOB)
- **External Tables:**
  - Format "informix" is not supported with BLOB, CLOB, BYTE or TEXT types
  - Must use Deluxe mode to load BYTE/TEXT types
- **HPL:**
  - Cannot use no-conversion jobs
  - Must use Deluxe mode to load BLOB, CLOB, BYTE or TEXT

Combine Methods

Utility	Use
dbexport/dbimport	Create initial schema Migrate smaller tables – exclude large
SQL Unload/Load	Small tables
dbload	Load only Target database uses logging Delimited & fixed format files Small-medium tables
External Tables	Medium-large tables Superior unload performance Flexible
HPL	Medium-large tables Superior load performance Flexible

Use PDQ

(not available in all Informix versions)

- Set PDQPRIORITY
- Parallel reads when unloading
- Improves utility performance
- Improves index builds
- Also consider PSORT_NPROCS
- Set DS_* values temporarily
- Be careful to avoid gating

Environment

- Source environment should match target
  - DBDATE
  - DBDELIMITER
  - DELIMIDENT
  - Others…

```
1204: Invalid year in date
 847: Error in load file row 1.
```

Performance Comparisons

- Timings performed against a **non-fragmented** table with 14 million rows, 440 MB
- Linux, IDS 12.10.FC5
- All loads performed into a RAW table with no indexes

Performance Comparisons

Unload

	No PDQ (seconds)	PDQPRIORITY 50 (seconds)
SQL Unload	26	22
External Table	28	11
External Table – 3 files	27	11
External Table – informix format	3	4

Performance Comparisons

Load

	No PDQ (seconds)	PDQPRIORITY 50 (seconds)
SQL Load	149	167
DBLOAD	145	150
External Table – Deluxe	156	110
External Table - Express	36	20
External Table – 3 files	33	16
External Table – informix format	23	17

Performance Comparisons

- Timings performed against a **fragmented** table with 16.5 million rows, 1.8 GB
- AIX 7.1, IDS 12.10.FC8
- All loads performed into a RAW table with no indexes

Unload

	No PDQ (seconds)	PDQPRIORITY 100 (seconds)
SQL Unload	136	131
External Table	58	54
External Table – 3 files	58	54
Ext. Table, informix format	12	12
Ext. Table, informix format, 3 files	12	12
HPL Delimited	103	103
HPL Delimited – 3 files	39	39
HPL no-conversion	22	22
HPL no-conversion – 3 files	20	20

Load

	No PDQ (seconds)	PDQPRIORITY 100 (seconds)
dbload	430	415
SQL Load	318	300
External Table – Deluxe	203	216
External Table – 3 files	202	211
Ext. Table, Deluxe, informix	163	195
Ext. Table, Deluxe, informix, 3 files	200	185
External Table – Express	162	181
External Table – 3 files	176	164
Ext. Table, Express, informix	150	126
Ext. Table, Express, informix, 3 files	139	127

Load

	No PDQ (seconds)	PDQPRIORITY 100 (seconds)
HPL, Delimited, Deluxe	108	108
HPL, Delimited, Deluxe – 3 files	100	102
HPL, Delimited, Express	52	52
HPL, Delimited, Express – 3 files	24	24
HPL, No-conversion, Express	18	19
HPL, No-conversion, Express – 3 files	19	18

Outage

- Need to have an outage to perform the migration
- Testing is essential before switching over
  - Time the migration
  - Tune the migration
  - Validate the migration
- Freeze schema changes

Incremental Migrations

- May not be possible to perform entire migration during an outage
- Copy static *tables* ahead of time
- Copy static *data* ahead of time
- Remove insert/update/delete permissions on *supposedly* static tables
- Triggers
  - Create triggers on tables to capture keys of records that have changed
  - Apply changes to the target system after consolidating

Changing Data

- Source data changing during capture will create problems
- Can the tables be locked during unload?
- Shut down applications?
- Implement triggers to capture changes
- Referential integrity may be compromised on target system
- Applying changes on target may require FK constraints be disabled until copy complete

Changing Data during Unload

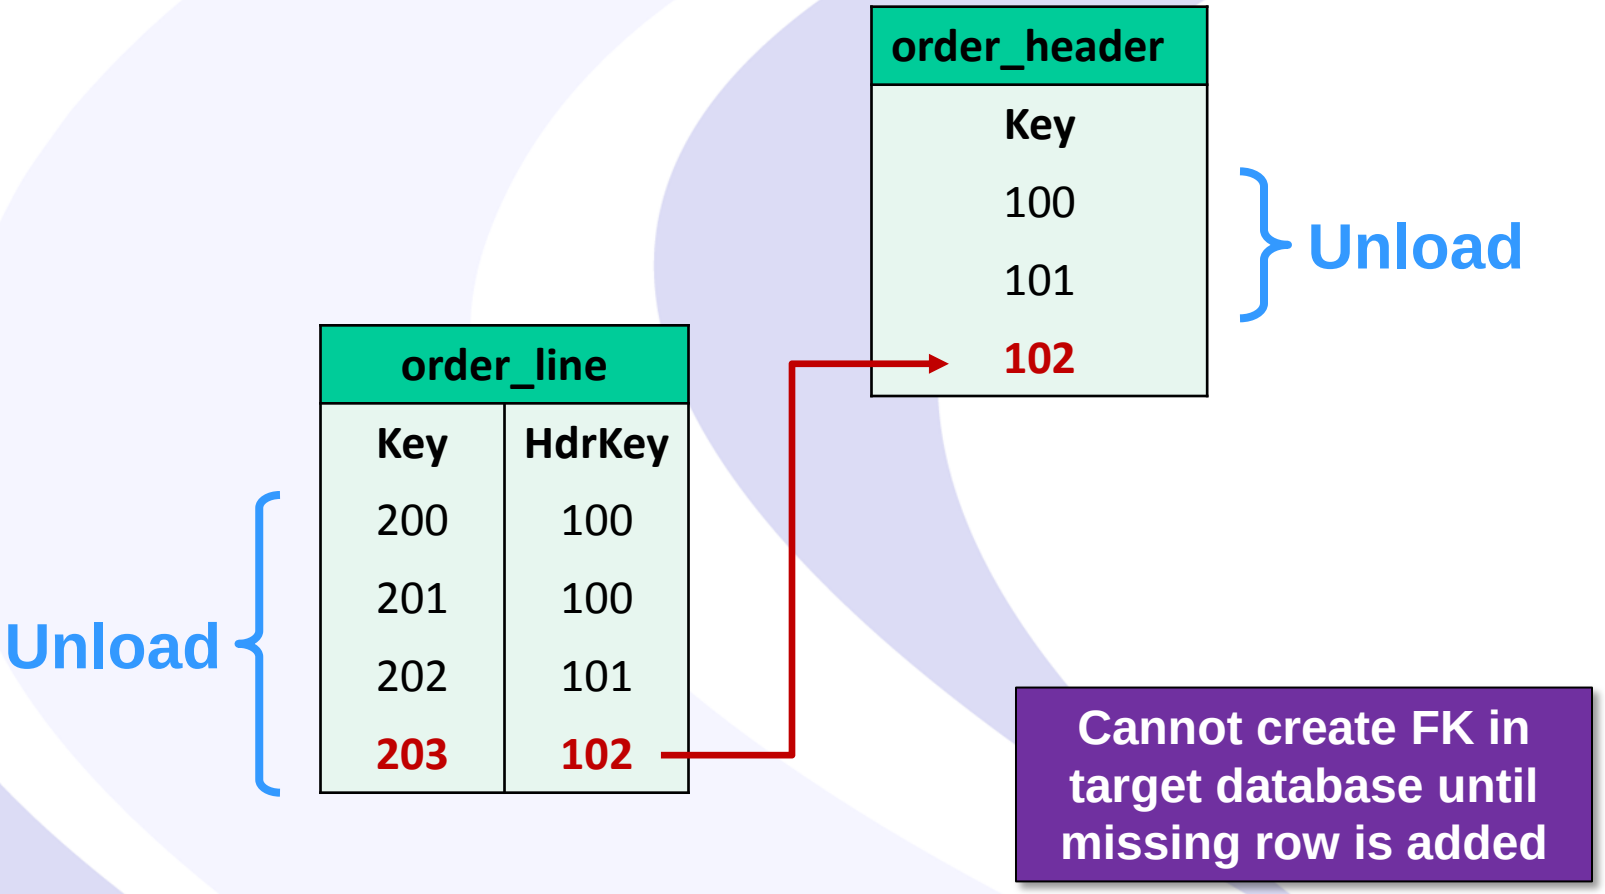

Enterprise Replication (ER)

- Use Informix Enterprise Replication to copy data over to target
- Target will remain synced with Source as changes are made
- Informix versions do not need to match!
- System architecture does not need to match!

Webcast on July 18th

Staging Area

- Staging filesystem required for unload and load files
- Disk write speed could be a bottleneck – avoid busy or slow disks
- NFS filesystem – bad configuration can cause corruption in unload files
- Larger the better
  - Avoid having to move, or zip files – takes time & CPU!
  - Avoid deleting unload files if possible

Informix Webcasts from the IBM Champions at Advanced DataTools

Managing Migrations With Informix ER

Thomas Beebe - Thursday, July 18, 2019 at 2:00pm EDT

Using Informix Enterprise Replication to move from an old server to a new one with minimal downtime

Registration and more information:
https://advancedatatools.com/Informix/NextWebcast.html

Informix Training
Updated for Informix 14.XX

Attend classes online on the web, or in person at our training center in Virginia. All you need is a web browser to connect to our WebEx training system and an SSH client (like Putty) to connect to our training lab for hands-on exercises. Each student uses an 8-core Linux server, with 16GB RAM, SSD drives with Informix 12, and several large databases for benchmark exercises.

- **March 11-14, 2019 - Advanced Informix Performance** — *Completed*
- **April 22-25, 2019 - Informix for Database Administrators** — *Completed*
  - This course is for new database administrators, programmers, and technical personnel who will be setting up, managing, and tuning IBM Informix databases.
- **October 7-10, 2019 - Informix for Database Administrators**
  - This course is for new database administrators, programmers, and technical support personnel who will be setting up, managing, and tuning IBM Informix databases.

- **More Information and Registration at:**
  http://www.advancedatatools.com/Training/InformixTraining.html

Advanced DataTools

Informix Training Servers

Each Student in class will have a server running Informix 12.10 with:

- 8 CPU Cores
- 16 GB RAM
- 1 SSD Disk
- 1- 4 Disks

IIUG World 2019

Register

https://www.iiug.org/conf/2019/iiug/register.php

North America IBM Informix Roadshow by Carlton Doe

- This series of free, one day, deep dive events is focused on the features of the recent IBM Informix v.14.10 release. A partial list of topics includes:
  - The new installation mechanism
  - New security enhancements including integrated backup encryption and changes to on-disk encryption!
  - High Availability enhancements
  - New SQL features and functions and application development enhancements
  - New replication functionality including asynchronous post commit triggers
  - Heterogeneous instance migration with codeset conversions
  - The Informix HQ graphical monitoring tool and more

- Please visit: http://tinyurl.com/mjmptfo

Questions?

Send follow-up questions to
info@advancedatatools.com

Informix Support and Training from the Informix Champions!

Advanced DataTools is an Advanced Level IBM Informix Data Management Partner, and has been an authorized Informix partner since 1993. We have a long-term relationship with IBM, we have priority access to high-level support staff, technical information, and Beta programs. Our team has been working with Informix since its inception, and includes 8 Senior Informix Database Consultants, 4 IBM Champions, 2 IIUG Director's Award winners, and an IBM Gold Consultant. We have Informix specialists Lester Knutsen and Art Kagel available to support your Informix performance tuning and monitoring requirements!

- ***Informix Remote DBA Support Monitoring***
- ***Informix Performance Tuning***
- ***Informix Training***
- ***Informix Consulting***
- ***Informix Development***

Free Informix Performance Tuning Webcast replays at:
http://advancedatatools.com/Informix/Webcasts.html
Email: info@advancedatatools.com
Web: http://www.advancedatatools.com

Thank You

Advanced DataTools Corporation

For more information:

info@advancedatatools.com

http://www.advancedatatools.com