

Migration Techniques for Oracle from distributed systems to a System z Cloud

Presented by: Rick Murphy, Migration Technical Solutions Architect





Problem Statement

What's the problem?

Migrating a database would appear to be a straightforward process but when analyzed critically it can become a complex undertaking when one or more of the following conditions exist:

- The source database is large
- The migration window is small
- The source and target operating systems are dissimilar
What do I copy? Is it compatible?
- The source and target systems are geographically remote
How do I transport the data?



Methodology Considerations

Migration Strategies

Three migration strategies exist:



*Methodology
Considerations*

Staged migrations are ones where specific parts of the database are migrated prior to the final migration window, reducing the data needing to be moved during the final migration window (and hence the downtime). This strategy is used for databases where portions of the data do not change over time, e.g. Data Warehouse

Big Bang migrations are ones where the complete database is migrated within the allocated window. This type of migration is used for OLTP databases which typically require that the whole database is moved to ensure consistency.

Zero Downtime migrations are ones where a very small migration window exists (< 1 hour). This type of migration will almost always require the use of a replication methodology.

Data Encapsulation



Methodology Considerations

There are two data transfer choices:

Physical data transfer involves transferring the *operating system level objects* (e.g. volumes, file systems, files) that contain the source database onto the target system.

It should be noted that for cross platform migrations, Oracle only supports physical transfer for versions 10g and 11g using transportable tablespaces / databases

Logical data transfer involves the transferring of the *logical objects*, such as tables, that constitutes the source database onto the target system, via use of database utilities (e.g. Oracle export/import, DataPump, INSERT..SELECT..,)

Data Transportation Mechanisms

	Tape	SAN	NAS	Replicate	Network
Physical or Logical	P or L	P	P or L	L	L
Reorganize / Consolidate Database					
Hardware Compatibility					
Lack of System Locality					
Operating System Compatibility					
Physical Level Copying (entire database)					
Double Handling of Data					
Cost of Tools					
Complexity					
Availability of swing storage/server					
Long set-up time between migration runs					
Initial 'Point in time' creation issues					
Migration Downtime	Slow	Medium	Medium	Fastest	Fast
Project Duration	Slow	Fast	Medium	Slow	Fast

- Requires additional work to provide capability
- Prerequisite
- Increased effort
- Reduced throughput
- Potential Show stopper

And the winner is Logical + Network



Methodology Considerations

This methodology has several important advantages in comparison to other methodologies including:

1. The elimination of cross platform issues
2. The elimination of data double handling
3. The elimination of any hardware components (tape systems, staging areas etc.)
4. Being a logical migration methodology, a significant reduction in the amount of data needing to be migrated (no indexes, no free space etc.)
5. The ability to reorganize the database
6. The ability to go from any release to any release
7. Natural de-fragmentation due to rows being reloaded into target

Final Thoughts



*Methodology
Considerations*

First law of database migrations:

“Don't touch the source database”

"Just because you can, doesn't mean you should"



Methodology Considerations

If a client has 10g on the target, don't immediately assume Transportable tablespaces is the solution. If there is an available window, do a logical migration and take the opportunity to:

- do a clean target build
- change the extent sizes on those large tables
- reorg that large monolithic table into a partitioned table
- boost performance from the defragmented target tables
- redesign the filesystem layout for performance

The major issue with upgrading a source database to 10g to facilitate a CPTTS migration is you double the effort and risk:

- For many third party and custom apps, upgrading the PRODUCTION source database to 10g will almost always require an upgrade to the PRODUCTION source application and will require significant levels of retesting of PRODUCTION applications and other PRODUCTION dependant environments before the actual migration to the new server where you will do exactly the same set of work and testing!



What is IBM XenoBridge?

What is IBM XenoBridge

A graphic for the 'XenoBridge Overview' section. It features a stylized globe on the left, partially obscured by a circular ring of binary code (0s and 1s) that appears to be floating or orbiting. The text 'XenoBridge Overview' is written in a serif font, with 'XenoBridge' in a larger size and 'Overview' below it.

XenoBridge Overview

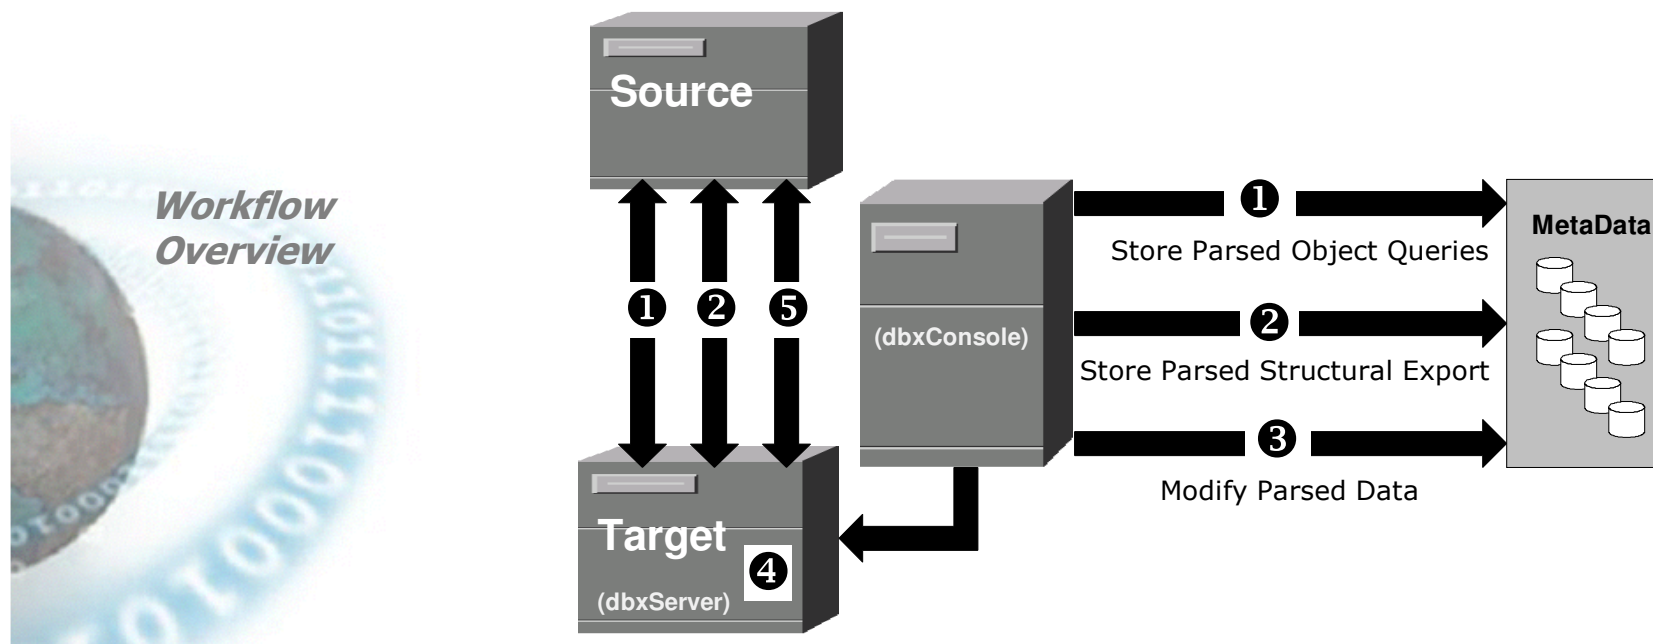
- IBM XenoBridge is a web based migration tool that utilizes advanced parallel processing techniques to migrate databases in a minimum of downtime (a little more than the time required to transfer the single largest table or partition)
- IBM XenoBridge utilizes vendor supported techniques to guarantee data integrity during the migration process
- Combined with an intelligent migration engine and multiple streams, the throughput obtained by IBM XenoBridge is up to 30 - 40 times that of export / import (average 6GB data / hour vs. average 200GB data / hour) *
- Data is transferred across the network meaning that the location of the source and target is transparent (state to state or country to country)
- IBM XenoBridge is fully featured and can not only move the rows in a database rapidly but also builds the target database based on the objects contained in the source database

* At Turkcell, XB moved a 9TB database in 6 hours!



Workflow Overview

Five steps to Migration Nirvana



1 Object queries are run against the source server and are stored in a metadata repository

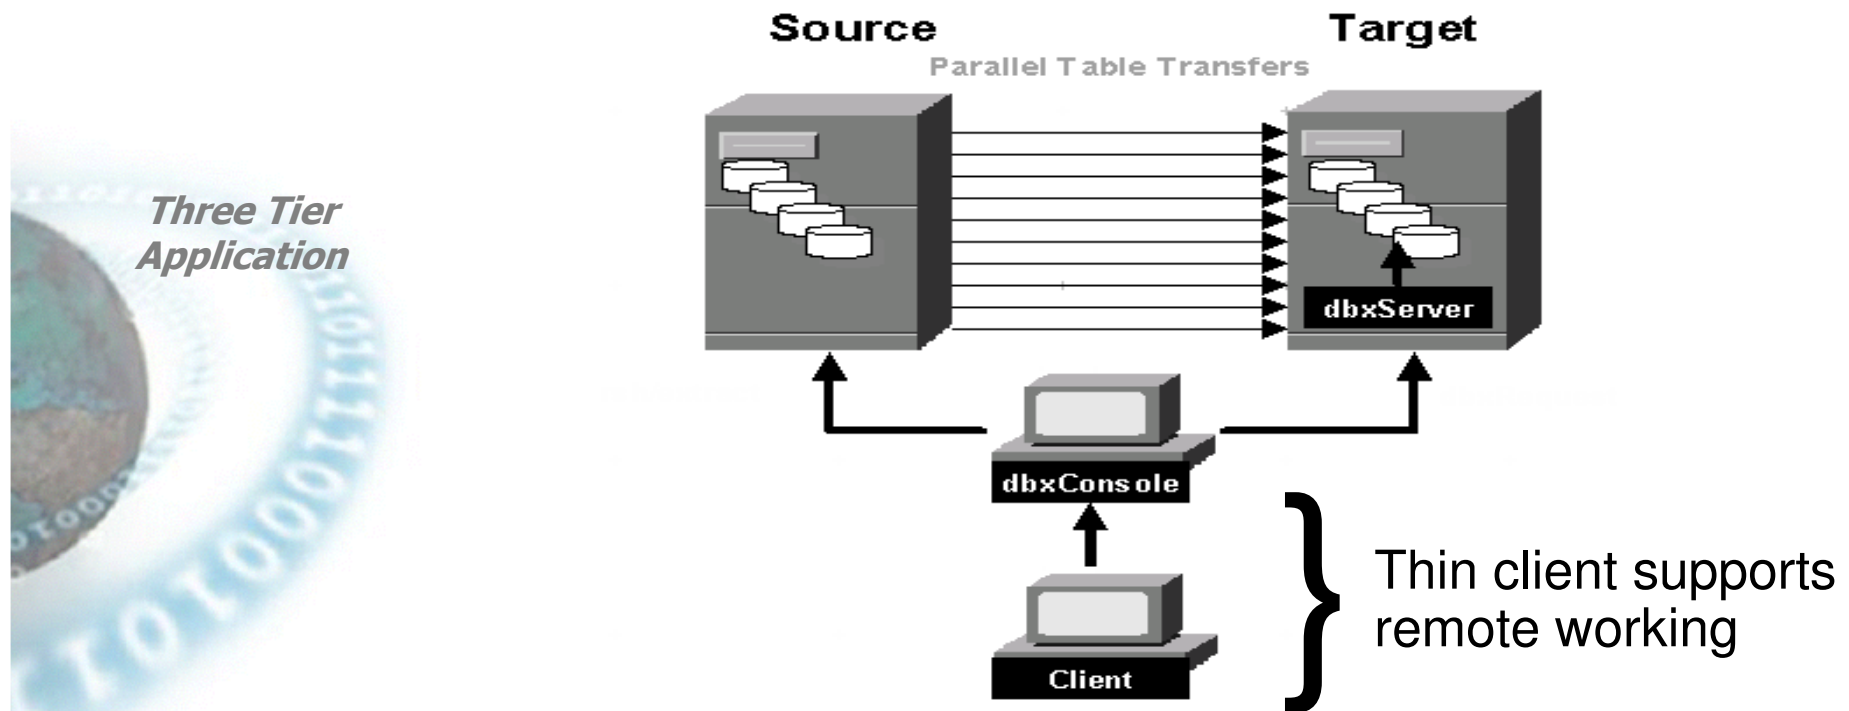
2 A structural export is run against the source server and specific DDL including create tablespace / user / table / index / constraint and grants are stored in the metadata repository

3 The DDL in the metadata repository can optionally be modified in specific ways to optimize, for example, storage allocation/location

4 The physical then logical database are created on the target and then validated to ensure all objects exist in the target database and that they have the same status

5 With a perfect “no-rows” database on the target server, the database rows are moved via the parallel migration engine

Component Overview



XenoBridge is a three tier application using a desktop based web browser which supports remote working

DbxConsole contains 99% of the 'code' and is responsible for extracting the source schema, building and validating the target database build, preparing the target database for a migration and initiating the migration

DbxConsole does not take part in the actual row movement

DbxServer is responsible for moving the rows in the database and validating that each table/table partition has moved successfully

Data is "pulled" from the source to the target

DbxServer moves multiple tables/table partitions concurrently

No software is installed on the source servers to facilitate the XenoBridge migration

Migrations – Some of the 'little' things that count

Target Database is pre-built prior to the migration

An empty target database including tables, indexes, stored procedures, grants, triggers etc is pre-built prior to moving the rows. By not having to create these objects as part of the migration itself significant time is saved

Move all or part of the database

Optionally, via inclusion/exclusion lists

Tables are moved in parallel

By moving tables in parallel, 30, 40 or 50 together, XenoBridge can utilize all of the resources available in high performance servers and SAN disk systems

Virtualize large tables

XenoBridge provides the ability to virtualize very large tables so that sub-sets of the table can be moved in parallel



Migrations – Some of the 'little' things that count

Empty tables are not migrated

No matter what methodology you employ it takes time to move empty tables. An empty table might typically take 3 seconds to migrate. If a database has 5000 empty that's 15,000 seconds or 4 hours of processing time saved (e.g. Oracle eBusiness Suite)

Indexes are rebuilt independent of the table load process

Indexes and Constraints are not transferred - loading a database with indexes applied causes unnecessary and unneeded overhead. XenoBridge disables indexes prior to the migration, rebuilding indexes as each table migration completes. As with table loads, index rebuilds are done in parallel. Foreign Key constraint rebuilds that fail are automatically restarted.

XenoBridge automatically adds in NOLOGGING, NOVALIDATE and PARALLEL hints to ensure indexes and constraints are rebuilt as quickly a possible

A graphic on the left side of the slide. It features a stylized globe on the left, partially obscured by a circular ring of binary code (0s and 1s) that appears to be floating or orbiting. The text 'XenoBridge Overview' is overlaid on the graphic in a serif font.

*XenoBridge
Overview*

Migrations – Some of the 'little' things that count (cont)

A graphic for 'XenoBridge Overview' featuring a globe on the left and a circular path of binary code (0s and 1s) on the right, all in a light blue, semi-transparent style.

XenoBridge Overview

Table migrations and index rebuilds auto-restart

Should an individual table migration or index rebuild fail, the table migration scheduler or the index rebuild scheduler will automatically restart the migration/rebuild process. The migration DBA can be paged or emailed when a restart occurs

Provides complete support for complex objects

Many database, particularly Oracle, allow the creation of complex objects (materialized views, advanced queues) which can be problematic to migrate. XenoBridge provides specific features to allow these types of objects to be migrated easily. XenoBridge contains optimizers for specific object types such as tables that contain LONG or LOB column types.

Upgrades Oracle versions on the fly

Upgrade your old Oracle 7/8/9 databases to versions 10/11

Migration Validation



Migration Validation

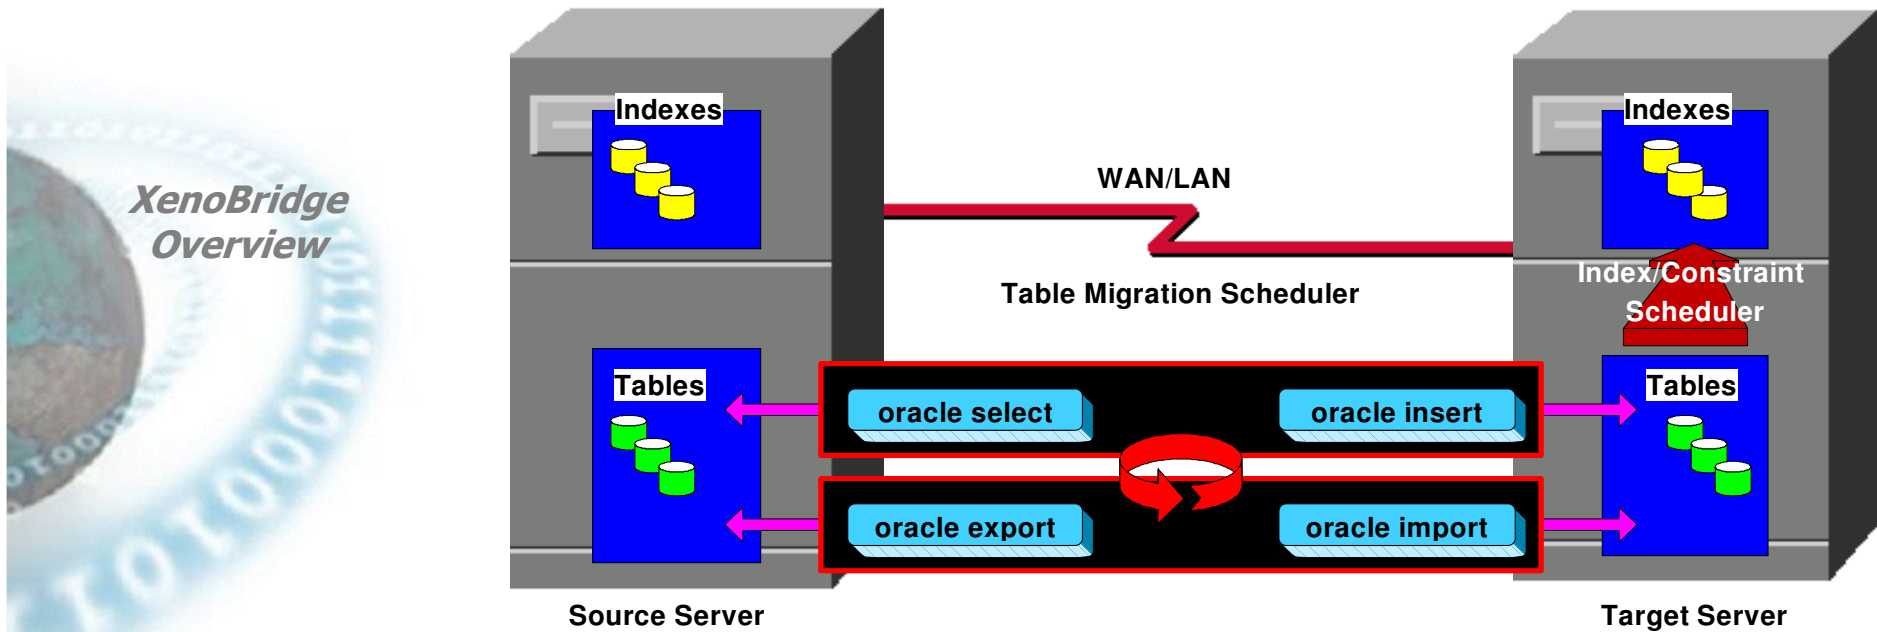
Prior to running a row migration and after building the database, XenoBridge compares every object in the source (tables, indexes, triggers, constraints, grants, stored procedures etc) with every object in the target.

During the row migration, the process that moves the rows scans the unload logs and the load logs, checking for database error messages (not enough space etc) and also matches the number of rows unloaded verses those that were loaded

Depending on the type of error encountered, XenoBridge will either automatically retry the row migration for that table or partition and/or send an email/page message

As a final row migration validation, a second row count is performed by comparing the row counts in the source with the row counts in the target for table or partition via a "SELECT COUNT (*). This second level validation ensures that all rows were correctly migrated

XenoBridge Topology - Oracle



99% of tables will move via INSERT with APPEND

Only LONG column tables move via exp/imp

Row movement is at partition level for partitioned tables

Large tables are virtualized via LTO to allow concurrent streams to be migrated

LONG column tables optimized via LCO to allow for faster row movement

Large LONG tables will be optimized with both LTO and LCO

Indexes and constraints are rebuilt during the row migration after each table has moved

Indexes and constraints used for the migration are optimized via PARALLEL hints, NOLOGGING and NOVALIDATE

Default parallelism is "degree 4" but indexes/constraints on the largest tables are optimized with higher degrees



Implementation Overview

Migration Project Steps

1 – Project Initiation

Project Kick-Off Meeting
 Project Review with Client
 Review database migration estimator script output
 Understand non-database migration phases
 Define database object relocation requirements
 Understand project constraints and timeline

2 – Database Build

Running of the logical database extract process
 Reworking of the Source database DDL where required
 Run the database creation DDL to create the database
 Rework the build DDL if required
 Validation of database creation

3 – Partial Migration test [any day of week]

Partial migration to validate setup & develop baseline
 Optional full migration not requiring database downtime
 Validation of migration results

4 - First Mock Migration Test

Partial migration to validate setup
 Prepare database for migration
 Complete migration of the production database
 Validation of migration results
 Post migration support during application testing by Client

5 – Optional Second Mock Migration Test [only if needed]

Partial migration to validate setup
 Prepare database for migration
 Complete migration of the production database
 Validation of migration results
 Post migration support during application testing by Client

6 – Go Live Migration Test

Partial migration to validate setup
 Prepare database for migration
 Complete migration of the production database
 Validation of migration results
 Post migration support during application testing by Client

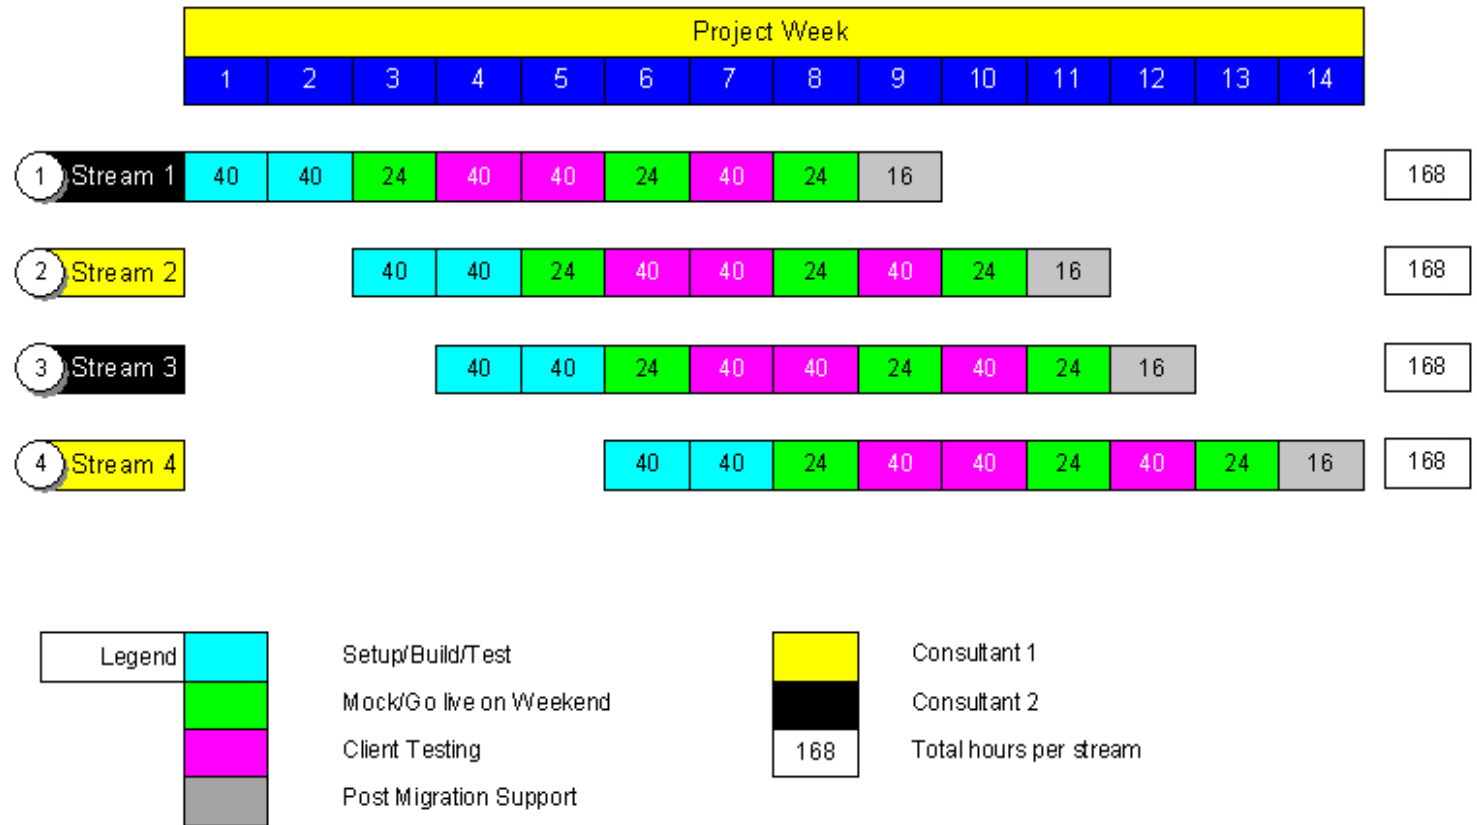
Only one Project Initiation is required
 Other steps are done per production database
 No mock migrations needed for development database
 The migration of a development/QA database is typically done just prior to its respective production database

Between each mock migration, the migrated database is available for client testing
 Other peripheral steps will exist if source code needs to be migrated and other infrastructure elements need to be reinstalled/upgraded/reconfigured

Implementation Example

This is an example of a multi-stream project utilizing two consultants

The feature to note is that the streams are delivered concurrently and that while the client is testing a stream, the consultant responsible for that stream can start work on another stream






Pre-Sales Analysis

To Each Their Own


*Pre-Sales
Analysis*



TCO - Total Cost of Ownership

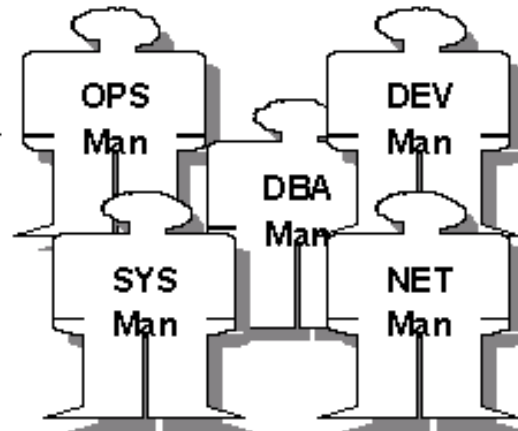
ROI - Return on Investment

To Each Their Own



TCO - Total Cost of Ownership

ROI - Return on Investment



TCO - Technically Complete Offering

ROI - Rapid Operational Integration

Xenobridge's Implementation Advantages

*Pre-Sales
Analysis*

By leveraging the database migration capabilities of XenoBridge, organizations will realize several major implementation benefits in comparison to other options:

Reduced Cost

- Customers can significantly reduce database licensing costs
- With every aspect of the migration life cycle already fully functional, the overall time taken to implement the migration project is significantly reduced

Reduced Risk

- The risks involved in migrating a database are reduced significantly and therefore the chances of meeting deadlines, staying under budget and ultimately being successful are greatly increased
- It's our experience that customers' fears about migration issues don't eventuate

Xenobridge's Implementation Advantages

Pre-Sales Analysis

Reduced Complexity

- XenoBridge has been designed from the ground as a fully featured database migration product, but it is not an “end user” tool
- Therefore XenoBridge is always provided as part of a packaged services engagement, freeing up your valuable technical staff
- The creation and tuning of large and complex databases requires not only in depth XB knowledge but also specialized Oracle migration knowledge, which we bring to the project

Your lives have been touched!!!!

GAP	(major USA retailer)
Paypal	(worlds largest online payment clearing house)
Western Digital	(world leader in disk drive technology)
Tengelmann Germany	(worlds 4th largest supermarket retailer)
Allianz Germany	(world's 2nd largest insurance and financial services organization)
Deutsche telekom Germany	(europe's largest telecom provider)
Opel Germany	(large european car maker)
Air Bus	(leading Aircraft manufacturer)
Disney	(major USA entertainment company)
Wells' Dairy	(largest family-owned and operated ice cream manufacturer in the US)



Any
Questions?

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 6 (logo), AS/400, Active Memory, BladeCenter, Blue Gene, CacheFlow, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, AIX 5L, Chiphopper, Chipkill, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Purpose File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER7, pureScale, System i, System p, System p5, System Storage, System z, Tivoli Enterprise, TME 10, TurboCore, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org. UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.

Intel, Itanium, Pentium are registered trademarks and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECcapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.

AltiVec is a trademark of Freescale Semiconductor, Inc.

Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.

InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.


Other company, product and service names may be trademarks or service marks of others.

Revised February 9, 2010



Backup Slides

What is not a Critical Consideration?

- 
- Source and target OS types/versions and source and target database versions are not critical factors
 - Size does not equal complexity or risk
 - Size is only a consideration relative to the downtime window
 - Moving databases supporting packaged applications (Oracle EBS,SAP,PeopleSoft, BAAN, Siebel, JD Edwards) does not increase risk but potentially only effort
 - Moving multiple databases within the same migration window is not a critical factor

What is a Critical Consideration?

The following are critical considerations but these do not increase risk, only duration:

- Large databases where we have limited access to the source databases for testing
- Migrations involving WANS
- Environments that contain replicats or failover copies of the database being moved or contain objects that have a master/master or master/slave relationship to other databases
- Migrations where the size verses downtime equation is challenging
- Migrations where the speed/reliability of the the source hardware is compromised
- Migrations where the database contains complex logical object relationships (inter and intra)



What is a Critical Consideration?



Critical Considerations

The following are critical considerations that increase duration and risk:

- Minimal downtime migrations that require the use replication software such as Golden Gate , SharePlex or IBM CDC
- Heterogeneous migrations from and to dissimilar databases



Network Performance

Typical Scenario

Network Performance

If we assume (which is almost always the case) that there is a single 1Gbit LAN connection between the source and target, a well configured target SAN (ie the target SAN is not a bottleneck) and a private network between the source and target:

- Then the THEORETICAL maximum throughput is 125MB/s or 439GB per hour. Of course we never get to this maximum but with tuning we can see up to 300gb per hour on a single link.
- Typically, Xenobridge can move 150gb to 300gb per hour of DATA with a minimum of tuning, based on the configuration of the network, the SAN and the number of links as described above. The bottom line is the more we tune the TCP/IP stack on the source and target servers and the more we tune XB, then the faster we will go.
- Which also brings into play the concept of DATA. Historically 50% of a database is non-data (indexes/free space/temp etc etc)... so a 300GB physical database , means XB only needs to move 150GB of data, which can be done within 1-2 hours.

Optimum Scenario - Deutsche Telekom

Network Performance



- Eliminated throughput bottlenecks prior to migration with 2 weeks tuning upfront on the target SAN and the network.
- With Four 1gb links we were able to drive each link at 80MB of data per second or 320MB across the four links which comes to 1TB of DATA per hour.
- XB can drive the infrastructure to maximum capacity, so the more care taken in tuning the infrastructure, the faster it will go

Legacy Scenario

Network Performance



- Potentially version 6, 7 and 8 databases may be on older hardware that only supports 100Mbit/s network cards
- This means that the practical throughput is one tenth of a 1Gbit link or around 12gb to 30gb of DATA per hours
- To assist migration speed in these legacy environments, XenoBridge can include compression within the data stream so the effective network throughput can be 4-5 times the raw throughput, meaning we can drive the legacy network in the 50gb to 120gb per hour of DATA range



Example Projects

Sample Client List

Xenobridge has been used at major corporations around the world for their mission critical database migrations including:

- **Paypal** (worlds largest online payment clearing house)
- **Western Digital** (world leader in disk drive technology)
- **Well Fargo** (major USA bank)
- **StateStreet Bank** (major USA commercial bank)
- **NCR** (major USA technology company)
- **Disney** (major USA entertainment company)
- **GAP** (major USA retailer)
- **Adecco** (worlds largest human resource company)
- **Ercot** (major USA energy Broker)
- **New York Dept. of Corrections** (major USA Incarcerator)
- **Williams** (major USA energy supplier)
- **Deutsche telekom Germany** (Europe's largest telecom provider)
- **Tengelmann Germany** (worlds 4th largest supermarket retailer)
- **Opel Germany** (large European car maker)
- **Allianz Germany** (world's 2nd largest international insurance and financial services organization)
- **Fiat Italy** (large European car maker and 6th largest in the world)
- **Turkcel** (Turkey's largest cell phone company)
- **Cosmote** (Greece's s largest cell phone company)
- **Air Bus** (Leading Aircraft manufacturer)
- **Wells' Daily** (largest family-owned and operated ice cream manufacturer in the US)



Project Estimation

Database Discovery Scripts



*Estimation
Overview*

- Cornerstone of the estimation process
- Written in SQL language
- One single integrated script for capturing database statistics
- Ensures a rapid turnaround time from assessment to proposal
- Standard spreadsheet can be easily populated with output from the script
- Totally risk free - it doesn't modify/delete any customer data, and because the source is provided in SQL format, customers can ensure that no confidential information is provided
- Output can be compared to previous projects to allow us to quickly estimate scope and costs

Client Questionnaire



*Estimation
Overview*

- In conjunction with the database estimator script, we also have the client fill out a questionnaire that covers:
 - Current source hardware sizes
 - Current source database versions and locations
 - Proposed target hardware sizes
 - Proposed target database versions
 - Proposed target database reorg's
 - Inventory of all source server resident software
 - Downtime windows for each database