



Welcome to the Virtual IMS user group newsletter. The Virtual IMS user group at www.fundi.com/virtualims is an independently-operated vendor-neutral site run by and for the IMS user community.

Virtual IMS user group presentation

The latest webinar from the Virtual IMS user group was entitled, “IMS replication for high-availability”, and was presented by Scott Quillicy, CEO and Founder of SQData.

Scott informed the group that there were a number of considerations that organisation should make before going ahead. He suggested:

- Replication is a simple concept, but filled with nuances.
- Best practices are the key to success.
- For IMS replication, there are few viable options.
- Make sure you evaluate all options.

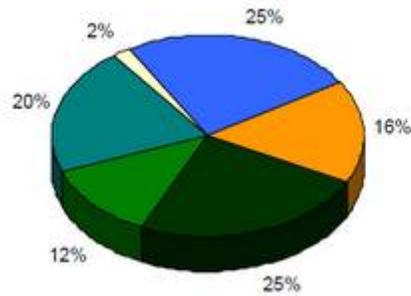


Figure 1: What CDC is being used for

- Call if you run into trouble.
- Planning:
 - Data movement requirements

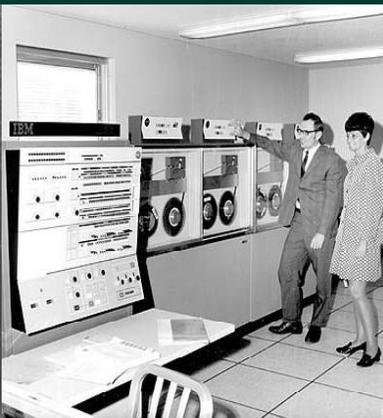
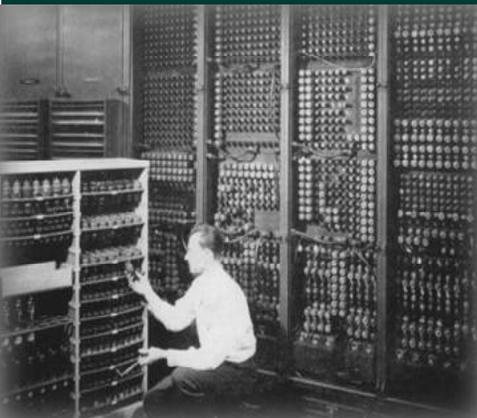
Figure 1 shows Scott’s experience of what Change Data Capture (CDC) is being used for at different organizations.

He told the group that High Availability was used by sites to minimize downtime, for disaster recovery, and for workload balancing.

Next Scott discussed replication best practices:

Contents:

Virtual IMS user group presentation	1
Meeting dates	5
Recent IMS articles	5
IMS news	5
About the Virtual IMS user group	5



- Latency requirements
- Changed data volumes
- Infrastructure requirements.

- Minimize the number of moving parts
- Have a recovery process
- Define conflict/exception processing criteria
- Have a method of verifying source/target consistency
- Have an operational monitoring strategy.

Any replication project has four primary components. These are source capture, data transport, target apply, and the all-important consistency monitoring – see Figure 2 below.

Simply put, replication can be divided into Active-Passive replication and Active-Active replication. Active-Passive is used for disaster recovery, Operational Data Store

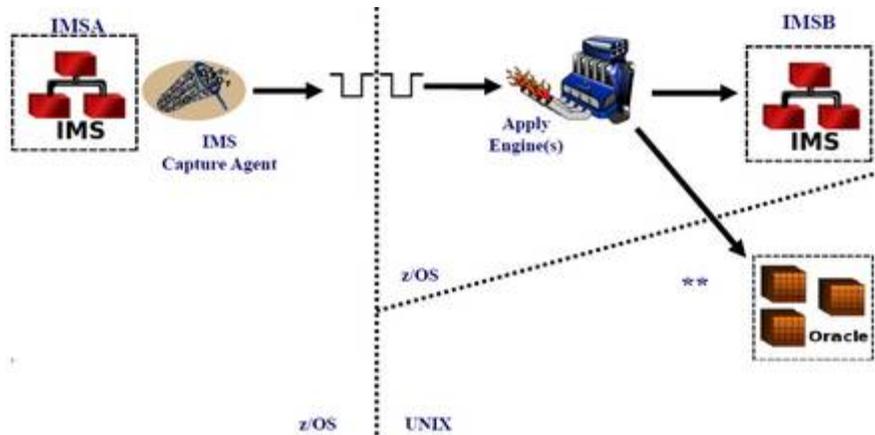


Figure 3: Active/Passive replication

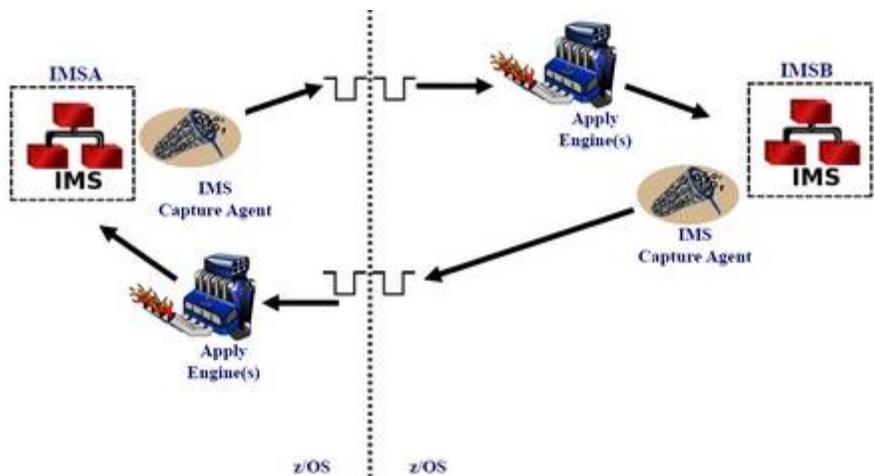


Figure 4: Active/Active replication

(ODS), and reporting, and is illustrated in Figure 3. Active-Active replication is used for continuous availability,

workload balancing, and disaster recovery, and is illustrated in Figure 4.

On top of this, it's possible to take some of the data captured and use it for downstream application processing.

Scott Quillicy went on to tell the user group that there were two primary method of IMS data capture – Data Capture Exit Routines and Log Based. The good things about Data Capture Exit Routines are: that it's near-real-time; there's

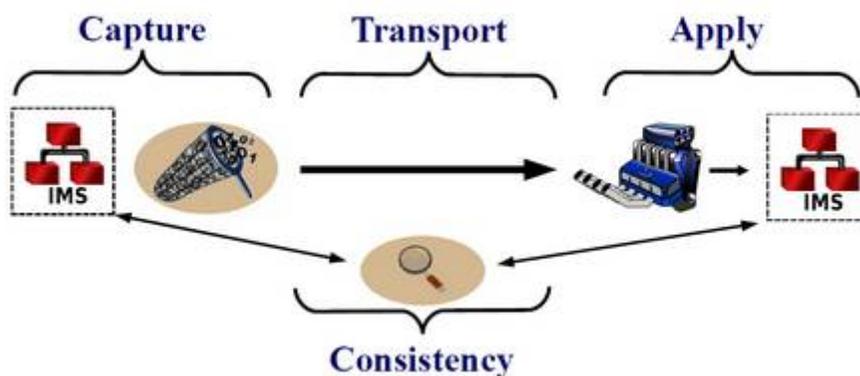


Figure 2: Primary components of replication

scalability; it can use MQ for persistent storage and transport; it doesn't require x'99' log records; and it executes in a dependent region as part of the transaction. This is illustrated in Figure 5 below.

The good things about Log Based data capture are: that it's near-real-time or asynchronous; it requires x'99' log records; it allows for recapture from SLDS; it's scalable; and it executes in a control region or in a separate address space. This is illustrated in Figure 6.

There are two primary methods for data transport – queue-based (eg MQ or Tibco), and native TCP/IP.

Queue-based data transport handles persistent storage in addition to transport; it's

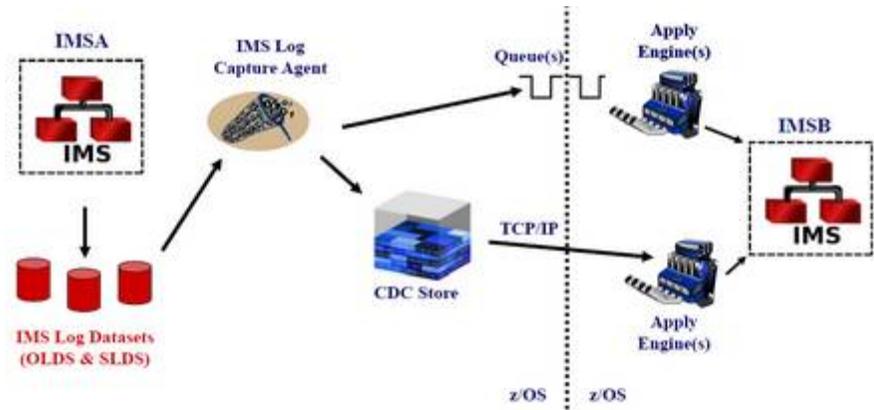


Figure 6: Log-based capture

resilient; it can handle high data volume on a continuous basis; and it operates independently of capture and apply.

Native TCP/IP transports are typically faster than queue-based; they can handle high data volumes on a continuous basis; requires separate storage for CDC data; resiliency must be built-in to

CDC storage; and operation is not always independent.

When looking at the Data Apply process, the data must be applied in order; it must be able to keep up; and it must be recoverable. For an Active/Active set-up, the process must avoid circular references; must maintain home affinity; and must provide conflict detection/resolution

In terms of Exception Processing: source exceptions include invalid data in a source change record, and are not common for relational CDC. With target exceptions, there can be issues with the target datastore (eg it being unavailable); conflicts; and application-specific logical errors.

A replication tool should allow for continued operation. It should include the ability to provide notification of

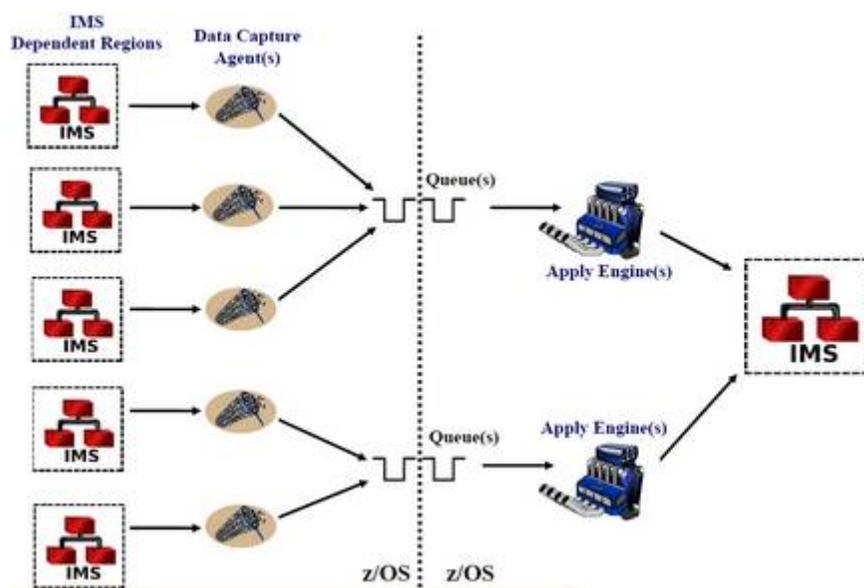


Figure 5: Data capture exit

exceptions; it should save exception records to a separate datastore; and some exceptions should force replication shutdown.

Scott recommended that conflict detection should be tied into consistency checking. In his experience, a conflict indicates a probable out-of-sync condition.

Replication conflicts occur when a segment wants to be inserted but already exists in the target; or a segment that does not exist in the target is trying to be updated; or a segment is being updated where the before image does not match the current target image; or a segment is being deleted that does not exist in the target; or any operation for a child segment that has no parent

But it's not all bad news, some conflicts can be compensated for. For example an insert for a segment that already exists can be turned into an update. Or an update for a segment that does not exist can be turned into an insert. And a delete for a segment that does not exist can be ignored.

Scott assured us that there were a number of options for resolving conflicts. The software could detect and compensate:

- COMPENSATE keyword on the target datastore

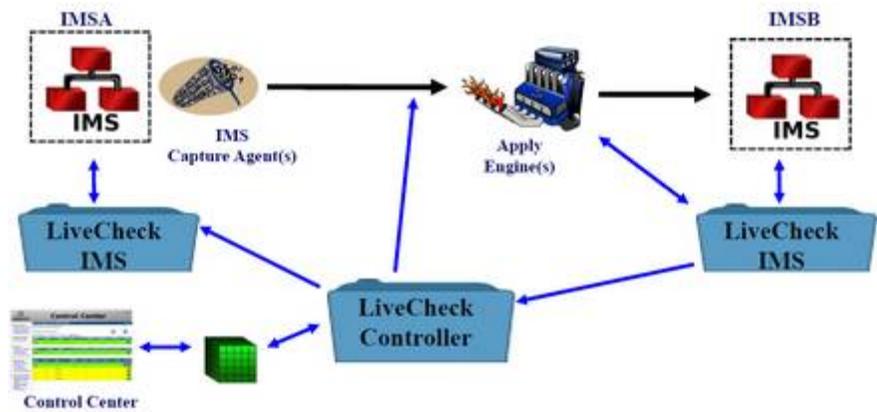


Figure 7: Verifying source/target consistency

- Compensate where allowed
- Write on-allowed compensations to exception datastore.

It could detect and compensate with a warning:

- COMPENSATE WITH WARNING keyword on the target datastore
- Compensate where allowed and issue a message for each occurrence
- Write non-allowed compensations to an exception datastore.

It could detect, not compensate, then continue running:

- Exception datastore specified for the target datastore
- Write conflicts to the exception datastore (DDLTO format).

It could detect, not compensate, and stop:

- No exception datastore specified for the target datastore
- Stops the apply engine.

The classic method to verify source/target consistency requires the source/target to be read-only:

- Unload source
- Unload target
- Transmit source unload to target (or *vice versa*), if required
- Compare source unload to target unload
- Not practical for fast moving databases.

A better approach is illustrated in Figure 7. It's intended for high-volume, 24x7 set-ups; and it checks the source and target integrity during active replication

Scott concluded with a summary repeating his initial considerations..

A copy of Scott's presentation is available for download from the Virtual IMS user group Web site at www.fundi.com/virtualims/presentations/IMSRepliationforHigh-AvailabilityAug11.pdf.

You can see and hear the whole user group meeting by downloading the WMV file from www.fundi.com/virtualims/presentations/2011-08-09meeting.wmv.

Meeting dates

The following meeting dates have been arranged for the Virtual IMS user group:

- 11 October 2011 – Rosemary Galvan, Principal Software Consultant – IMS, BMC. “Database Performance – Could Have, Should Have, Would Have”.
- 6 December – an IBM presenter, “IMS 12”.

Recent IMS articles

IBM's IMS Newsletter, July 2011 - Volume 1102 (<http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=c=2Fcom.ibm.imsnews.doc%2Fnewsletters%2Fv1102%2Fv1102.htm>) contains:

Enabling Mainframe Assets for SOA Using IMS Connect

and DataPower by Sean Hickman, Dave Holman, and Surajit Pal, Wellpoint

Transition IMS Disaster Recovery into Disaster Restart by Kelly Smith, Rocket Software

Business Process Integration for IMS Applications by Barbara Klein, IMS Business Programs Manager; Kevin Lo, IMS MFS Web Solutions Lead

IMS Enterprise Suite 1.1 SOAP Gateway: New Features by David Hanson, IMS Information Developer

IMS Security Support in IMS Enterprise Suite 1.1 SOAP Gateway by Jenny Hung and Quihong Sun, IMS SOA Solutions and IMS SOAP Gateway Software Engineer.

IMS news

IMS is one of IBM's 100 icons of progress. More information can be found at www.ibm.com/ibm100/us/en/icons/ibmims/.

CA Technologies has announced that CAACF2 r14 SP1 for z/OS and CA Top Secret r14 SP1 for z/OS have met the requirements of Common Criteria, an independent security certification recognized by governments in more than 26 countries, including the USA. CAACF2 and CA Top Secret provide access control for

z/OS resources across operating systems, subsystems, third-party software and databases, which includes externalized security controls for UNIX System Services (USS) as well as DB2, CICS, and IMS transaction and database management subsystems. More information can be found here. Full information can be found at www.ca.com/us/news/Press-Rel/2011/CA-Technologies-Mainframe-Security-Management-Products-Meeting-Rigorous-Requirements.aspx.

About the Virtual IMS user group

The Virtual IMS user group was established as a way for individuals using IBM's IMS hierarchical database and transaction processing systems to exchange information, learn new techniques, and advance their skills with the product.

The Web site at www.fundi.com/virtualims provides a central point for coordinating periodic meetings (which contain technically-oriented topics presented in a webinar format), and provides articles, discussions, links, and other resources of interest to IBM IMS practitioners. Anyone with an interest in IMS is welcome to join the Virtual IMS user group and share in the knowledge exchange.

To share ideas, and for further information, contact trevor@itech-ed.com.

The Virtual IMS user group is free to its members.