IBM Data Virtualization Manager for z/OS

Session Number: 22621

John Casey
Senior Solutions Advisor
Rocket Software
jcasey@rocketsoftware.com

Disclaimer

-
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be
-
-
-
-
-

Critical Business Issues

Data resides in silos and unfamiliar formats

Speed of business and new digital requirements forcing need for real-time information

There is simply too much data for older architectures

Traditional Data Integration

Challenges of moving data
(Traditional ETL)

- Risk to data security
- Data inconsistency
- Rigid, limits business agility
- High cost and latency

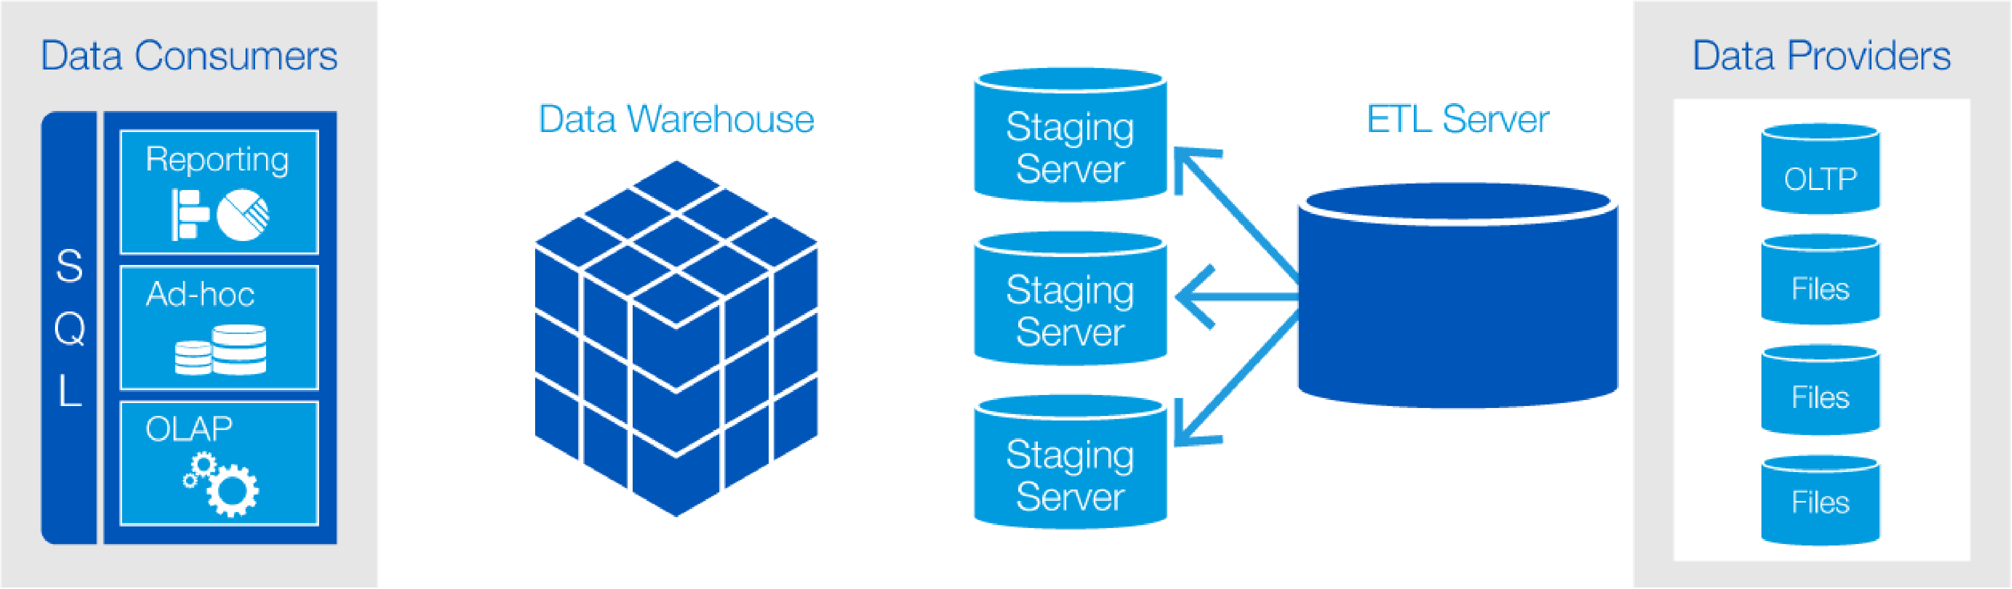

IBM Data Virtualization Manager for z/OS

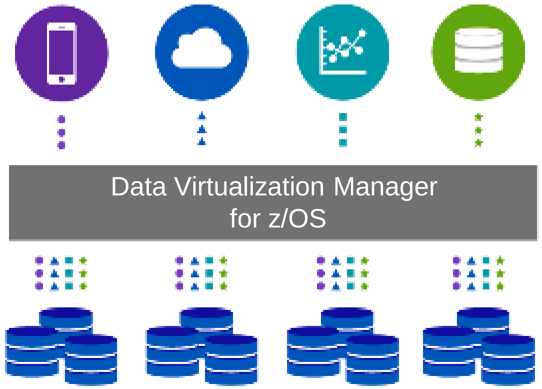

- IBM Data Virtualization Manager for z/OS
  - IBM® Data Virtualization Manager for z/OS® allows organizations to virtualize z/OS data with other enterprise data sources in real-time to provide comprehensive information that is readily consumable by analytics, digital applications and users.

- Value Proposition:
  - IBM Data Virtualization Manager for z/OS enhances IBM's premier z/OS data integration solution with new capabilities for creating real-time, virtual views of enterprise data - mainframe and non-mainframe. With Data Virtualization Manager, data remains securely in place, ready when needed in the right format, without the cost and complexity associated with moving data. Data Virtualization Manager unlocks the value of your mainframe data for real-time customer and business insights.

- Key Features:
  - Keeps data secure and in-place with real-time data virtualization on IBM Z
  - Support for a broad range of mainframe and non-mainframe data sources
  - Breadth of APIs and interface support for SQL, JSON, REST, SOAP and HTTP
  - z/OS resident optimization for improved performance and TCO
  - Abstraction layer to improve productivity and reduce reliance on mainframe skills

Data Virtualization Use Cases

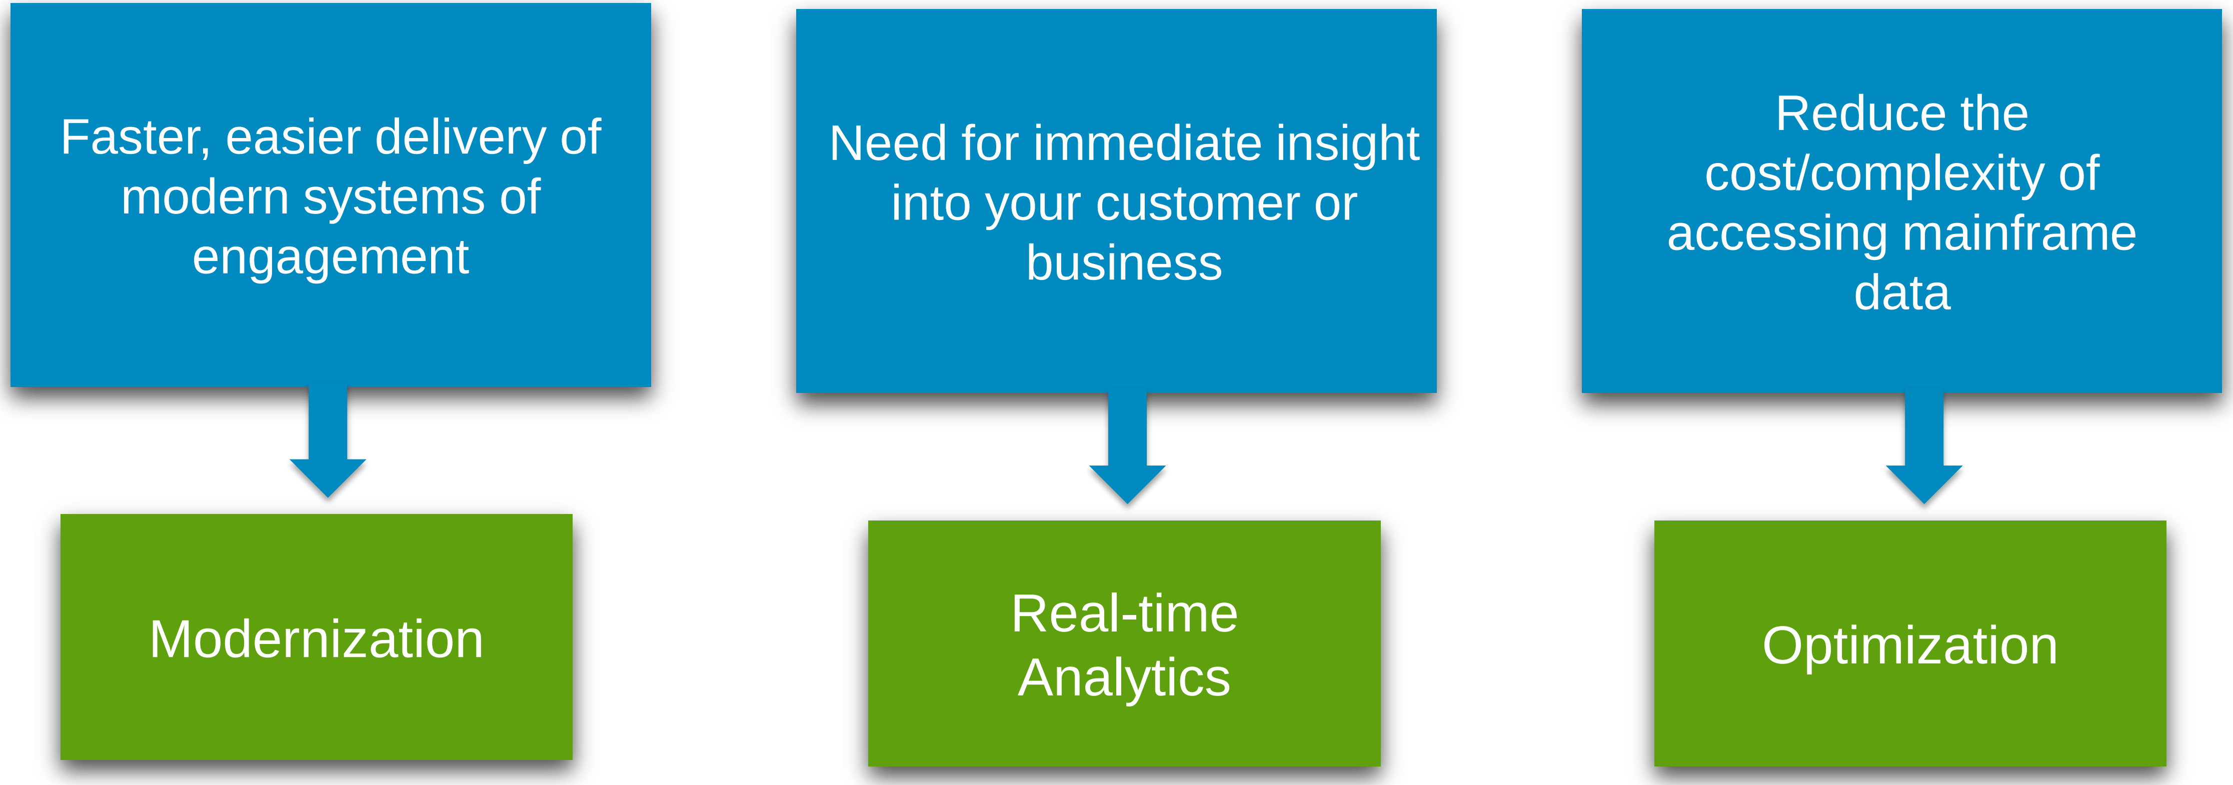

Faster, easier delivery of modern systems of engagement

Modernization

Need for immediate insight into your customer or business

Real-time Analytics

Reduce the cost/complexity of accessing mainframe data

Optimization

Data Virtualization Technology Accelerates Real Time Analytics

Universal Data Access

Data Virtualization for all enterprise data -
on and off mainframe

BI Solution

with Access, Virtualization and Visualization
and Data Preparation

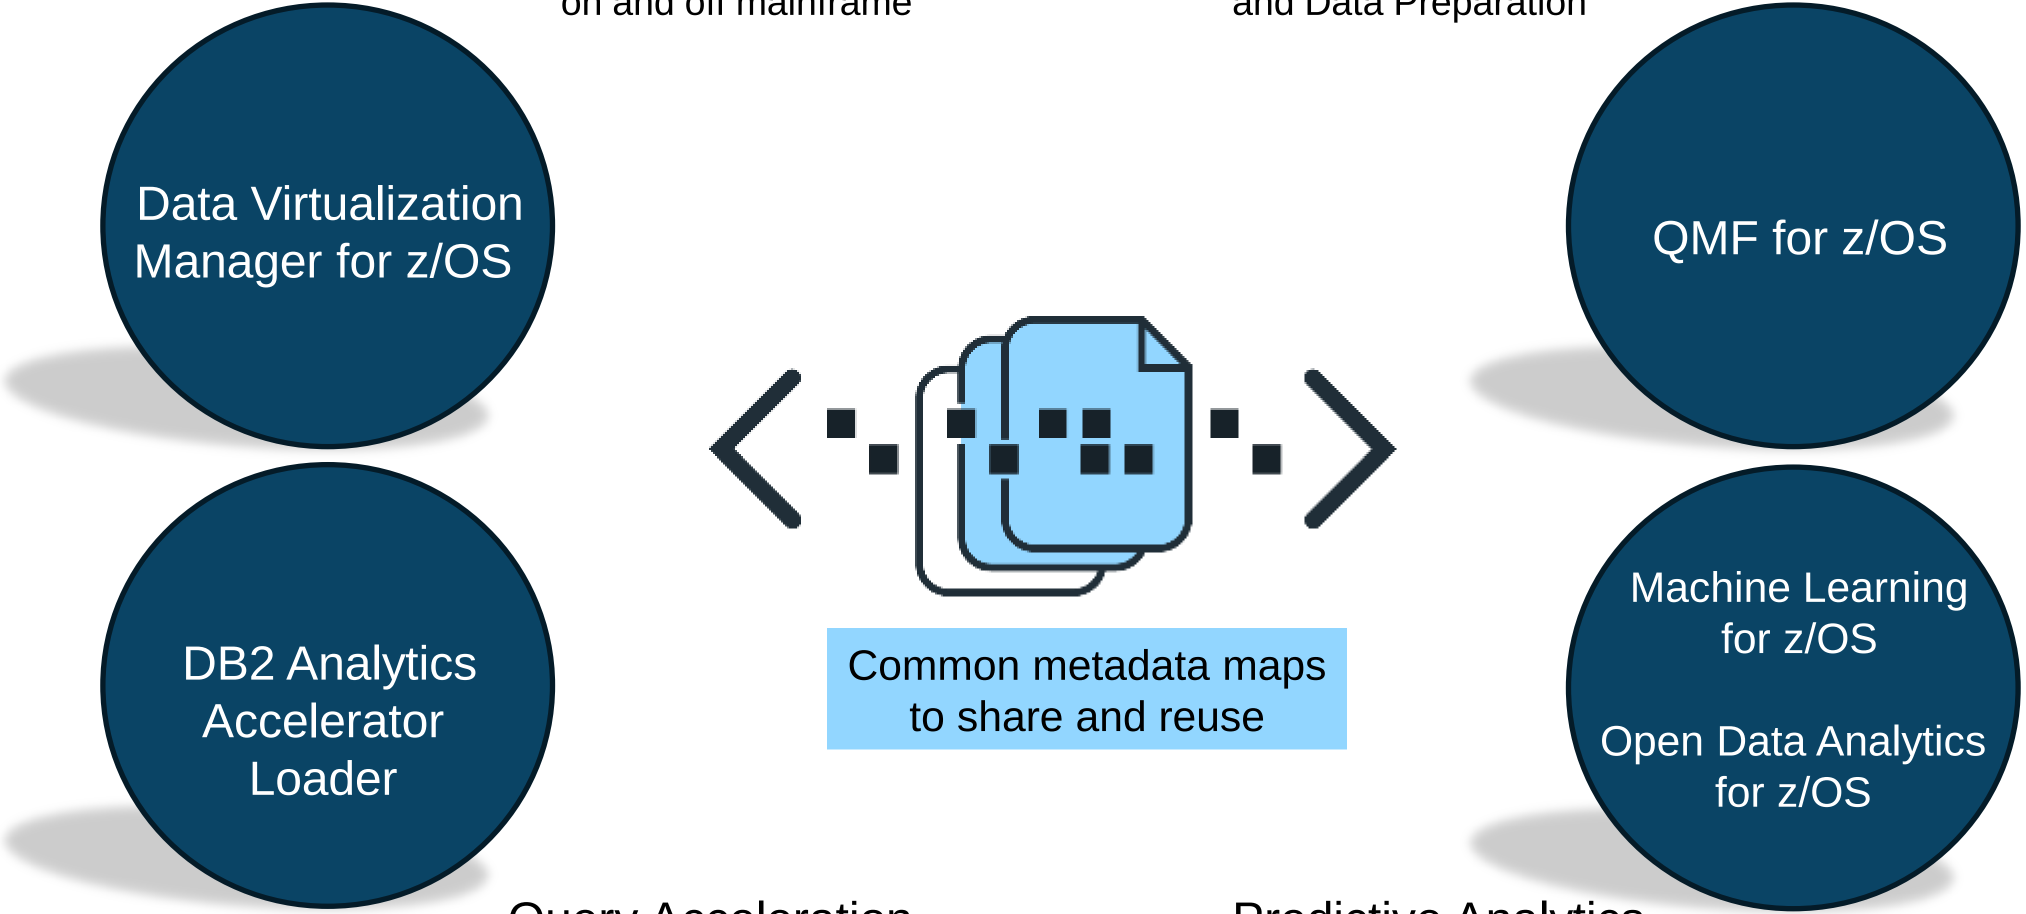

Data Virtualization Manager for z/OS

DB2 Analytics Accelerator Loader

QMF for z/OS

Machine Learning for z/OS

Open Data Analytics for z/OS

Common metadata maps to share and reuse

Query Acceleration

High-speed processing for complex Db2 queries

Predictive Analytics

Discovering patterns/meaning in data

IBM Data Virtualization Manager for z/OS

Virtualize z/OS data with other enterprise data sources in real-time to provide comprehensive information that is readily consumable by analytics, digital applications and users

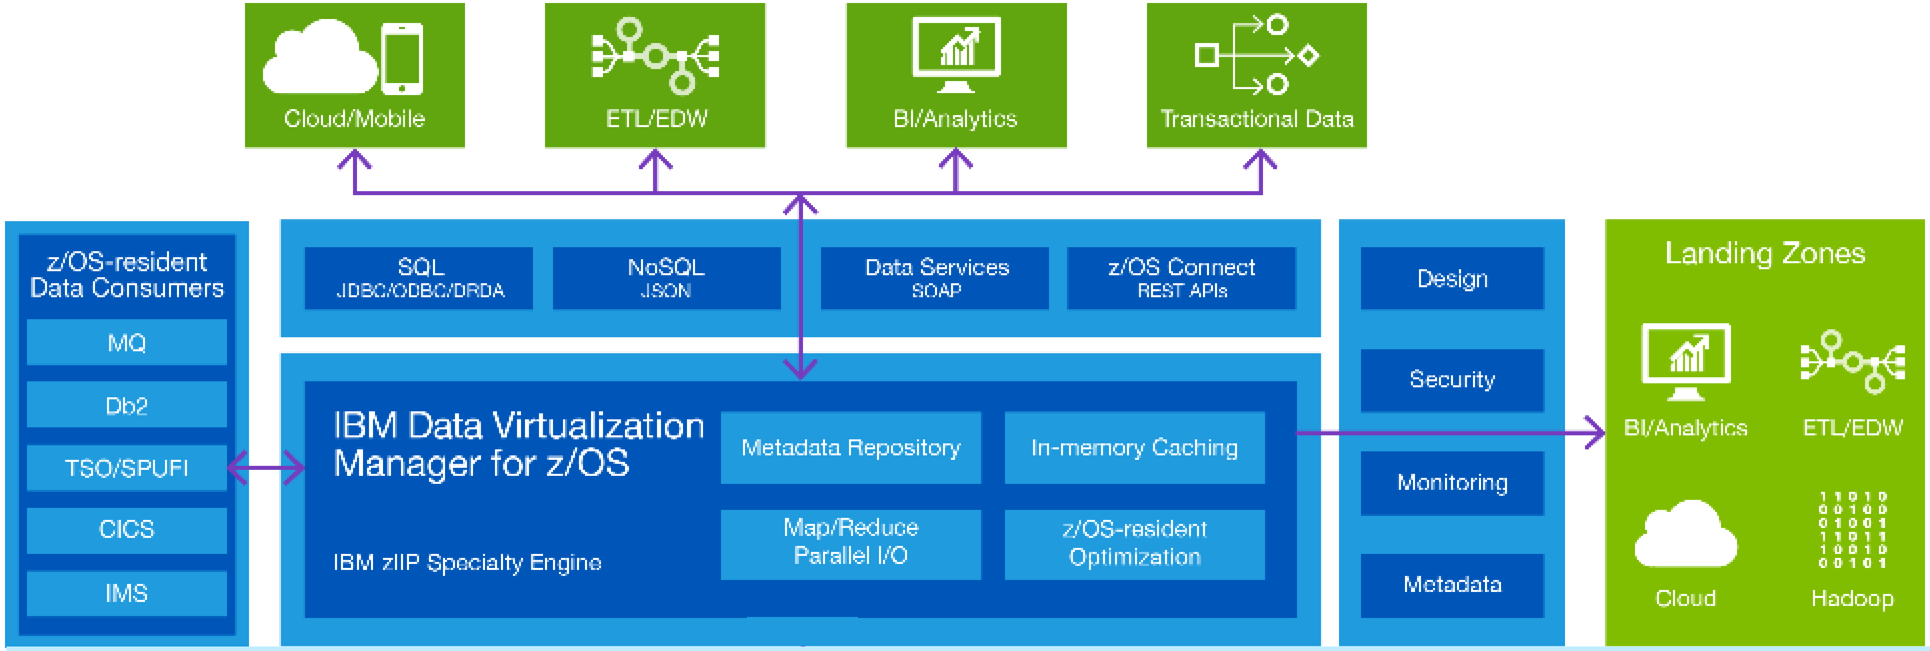

- zIIP eligible / avoid MLC costs
- Optimized for z14, z13, EC12
- Supports relational sources
- Rules engine

- Support any API:
  SQL, NoSQL, REST, SOAP, HTTP
- Python support
- Automated discovery though IBM ADDI

- Ships with industry mappings
- Exploits 64-bit storage
- Built-in MapReduce
- Built-in parallel I/O

New Feature: MQ Queues and Messages

- Ability to perform SQL-based analytics directly against MQ queues/messages
  - treats MQ queues/messages as a new data source
  - ability to filter on individual messages to include in result set
- Join any DVM data sources with MQ messages to access as virtual table
- Support for MQGET (reads) and MQPUT (inserts)
  - Non-invasive, virtual access
  - Eliminates the need to extract data from queue

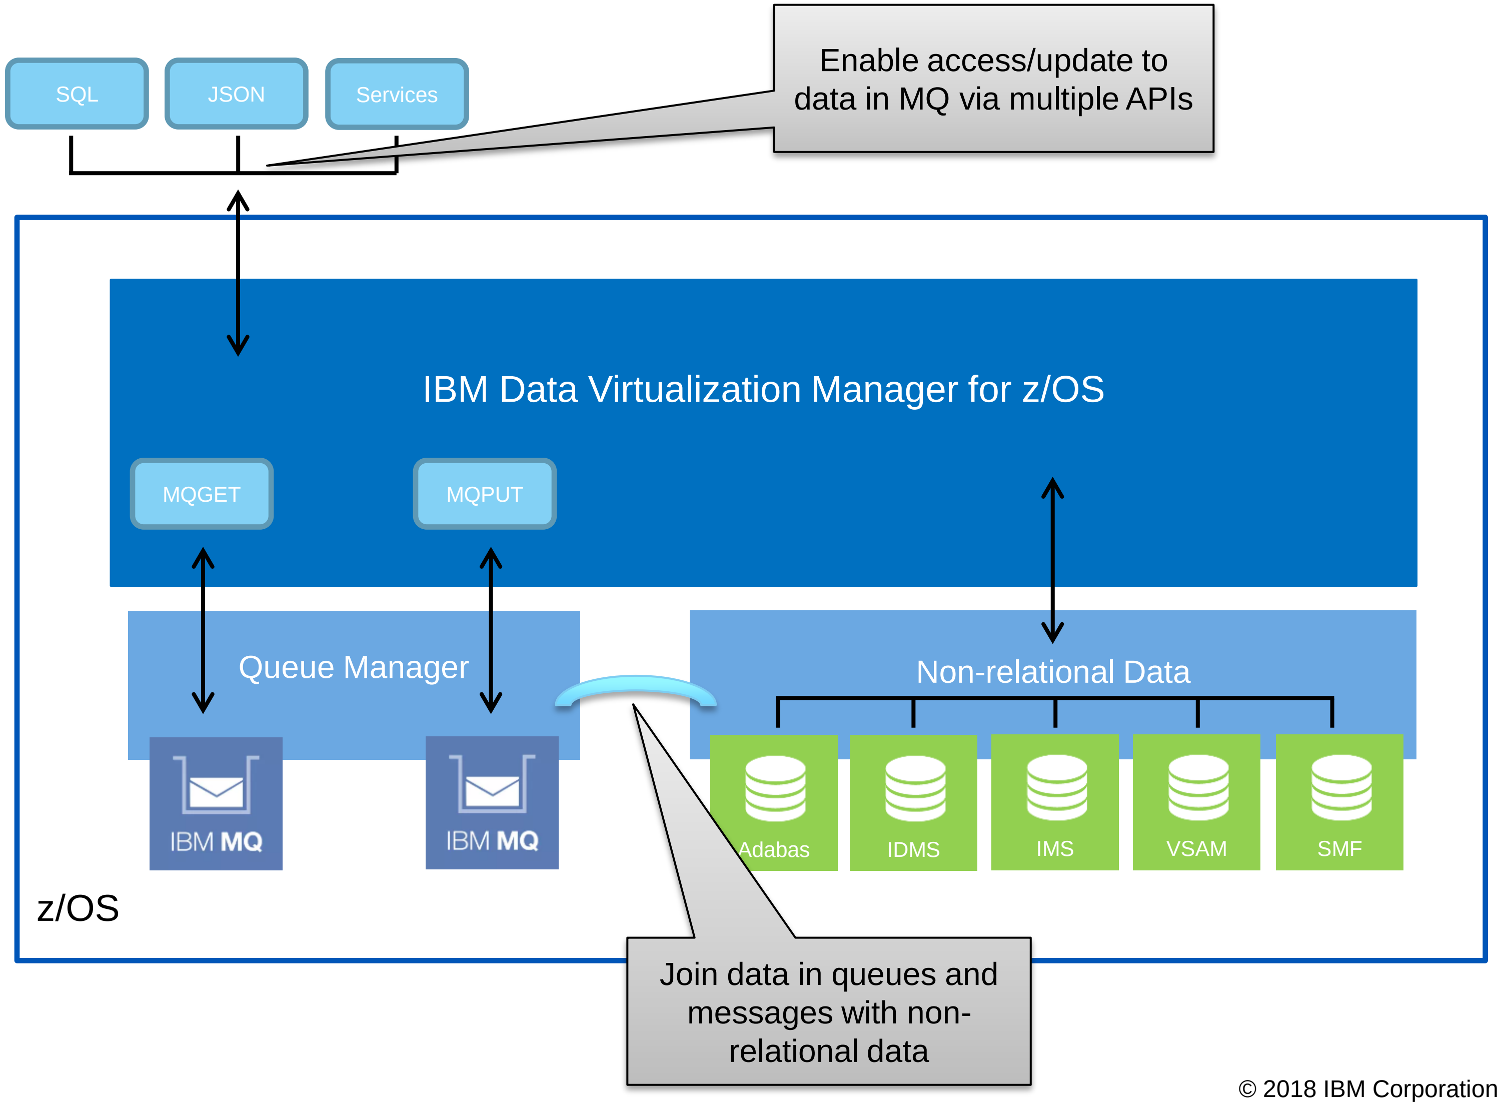

Enable access/update to data in MQ via multiple APIs

Join data in queues and messages with non-relational data

Virtual Parallel Data

- Allows multiple independent SQL requests to share a single result set
  - Front ends
    - IDAA Loader
    - IzODA - Spark Scala/Spark Java/Python DB API
    - QMF
    - JDBC/ODBC
    - Studio
    - Batch/DSSPUFI
  - Back ends
    - Sequential/Tape
    - VSAM/IAM
    - Adabas
    - zFS
    - Logstreams
    - MQ Series
    - IMS
    - IDMS

Virtual Parallel Data

- VPD requests are joined in a named group
  - ➢ Each request specifies the name of the group
  - ➢ The first request to arrive creates the group; I/O begins immediately.
  - ➢ Requests arriving within a time interval are placed into the same group.
  - ➢ The group is closed when the last request arrives, or the timeout is reached.
  - ➢ Once the group is closed, the buffer can be wrapped/reused.
  - ➢ The back-end I/O is done once, and buffered in one or more large, wrapping, 64-bit memory objects.
  - ➢ All group members share the buffered results.
  - ➢ Each VPD request can specify its own degree of parallelism using MRC.

Recent PoT Environment

Objective: Replace multi-step ETL process to replicate IDMS tables to Oracle

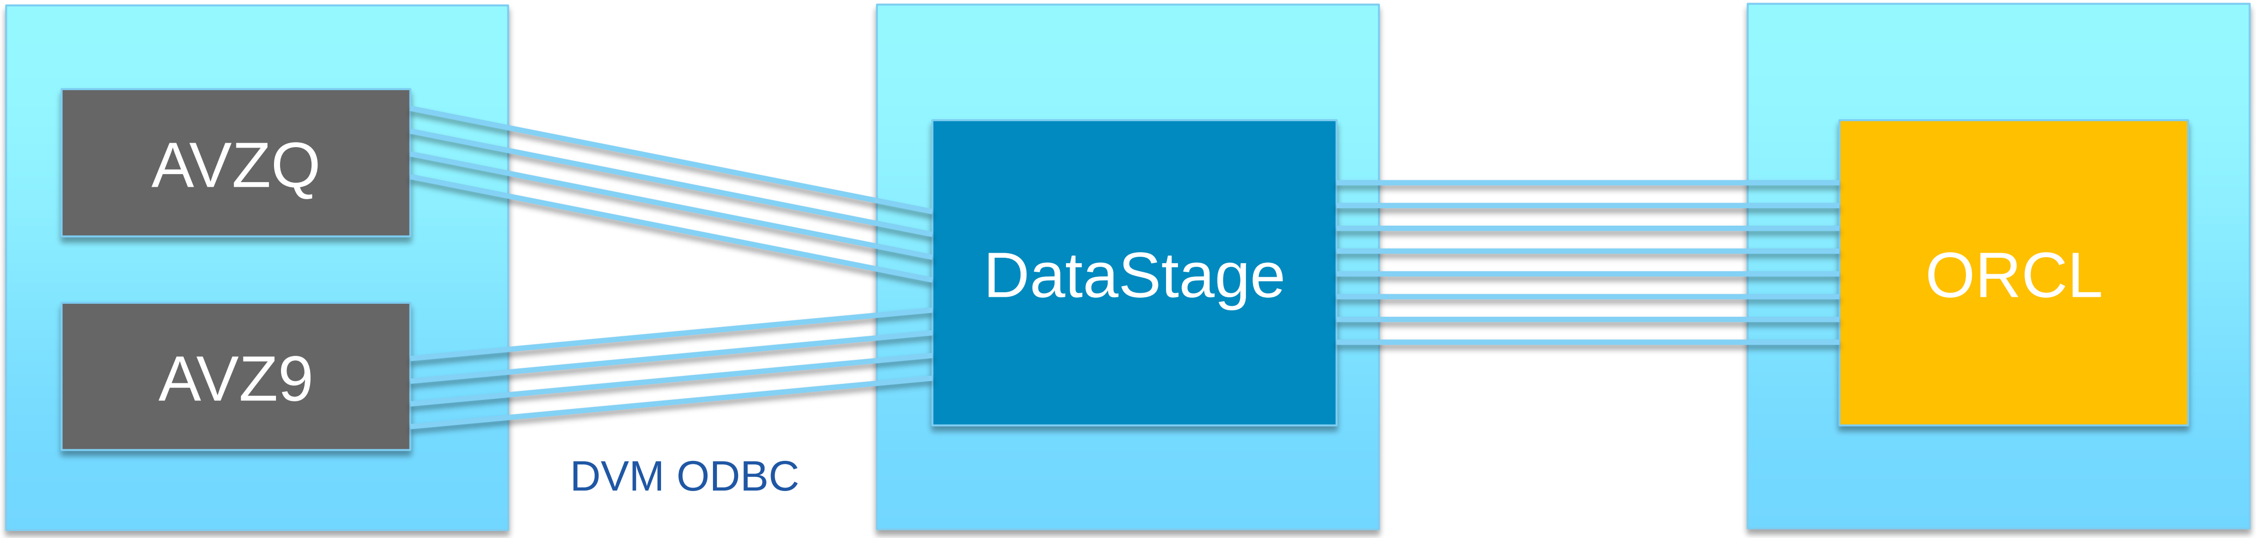

z13
- 768GB
- 7 GPP
- 6 zIIP

Windows VM
- 2 processors
- 8 GB

Linux VM
- 2 processors
- 8GB

ODBC Map Reduce

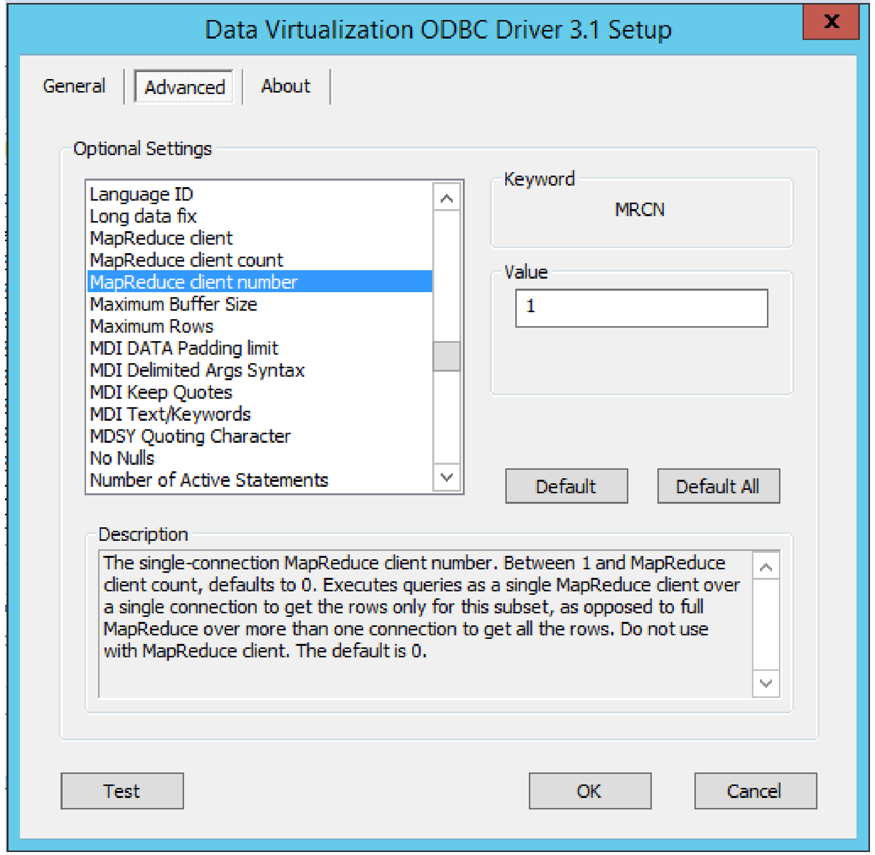

Optional Settings

Language ID
Long data fix
MapReduce client
MapReduce client count
MapReduce client number
Maximum Buffer Size
Maximum Rows
MDI DATA Padding limit
MDI Delimited Args Syntax
MDI Keep Quotes
MDI Text/Keywords
MDSY Quoting Character
No Nulls
Number of Active Statements

Keyword: MRCN

Value: 1

Description: The single-connection MapReduce client number. Between 1 and MapReduce client count, defaults to 0. Executes queries as a single MapReduce client over a single connection to get the rows only for this subset, as opposed to full MapReduce over more than one connection to get all the rows. Do not use with MapReduce client. The default is 0.

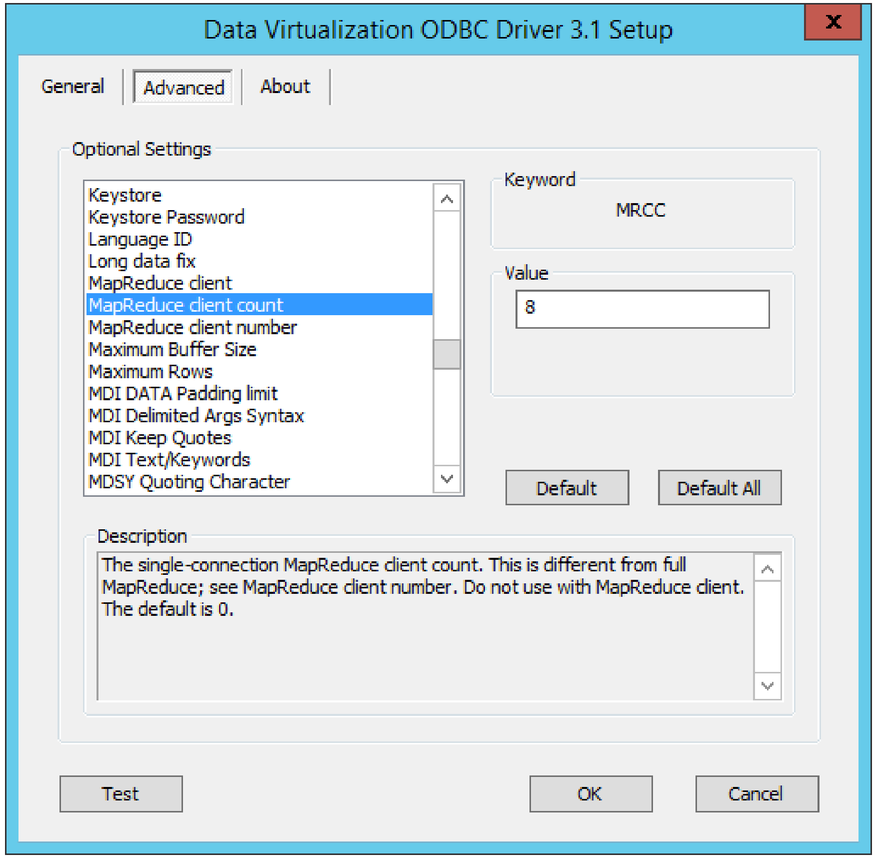

Optional Settings

Keystore
Keystore Password
Language ID
Long data fix
MapReduce client
MapReduce client count
MapReduce client number
Maximum Buffer Size
Maximum Rows
MDI DATA Padding limit
MDI Delimited Args Syntax
MDI Keep Quotes
MDI Text/Keywords
MDSY Quoting Character

Keyword: MRCC

Value: 8

Description: The single-connection MapReduce client count. This is different from full MapReduce; see MapReduce client number. Do not use with MapReduce client. The default is 0.

How does IBM Data Virtualization Manager perform?

Tests performed at the IBM Systems Benchmark Center, Poughkeepsie, NY in Nov. 2017 running on IBM z13 using 800GB of financial data - flat files, with a multitude of fields

zIIP exploitation

- 99% of data virtualization runs on zIIP

Parallelism

- Test Case 4 and 2 have same configuration
- With degree of parallelism at 8 elapsed time is reduced from 98.68 minutes to 17 minutes
- Furthermore, by adding 3 zIIPs, test case 8 shows even greater improvement bringing the elapsed time down to 13.83 minutes.
- With enough zIIPs it will not unusual for us to see 1000% improvement for elapsed times

zIIP engine exploitation

Row Labels	Sum of CPU Time	Sum of zIIP Time	Sum of IIPCP Time	Sum of zIIP NTime	%zIIP eligible
DVM 1	7099.03	5609.55	1389.58	5609.55	98.59%

Parallelism impact on elapsed time

Test Case	GPP's	Number of zIIP engines	Degree of parallelism	Elapse time in minutes	SMT
1	8	0	0	118.96	1
2	8	5	0	98.68	1
3	8	5	4	27.05	1
4	8	5	8	17.14	1
5	8	5	8	20.84	2
6	8	5	10	17	2
7	8	5	16	15.73	2
8	8	8	8	13.83	1
9	8	8	8	17.62	2
10	8	8	16	11.72	2

Data Virtualization for Mainframe API Enablement

z/OS Connect EE v3

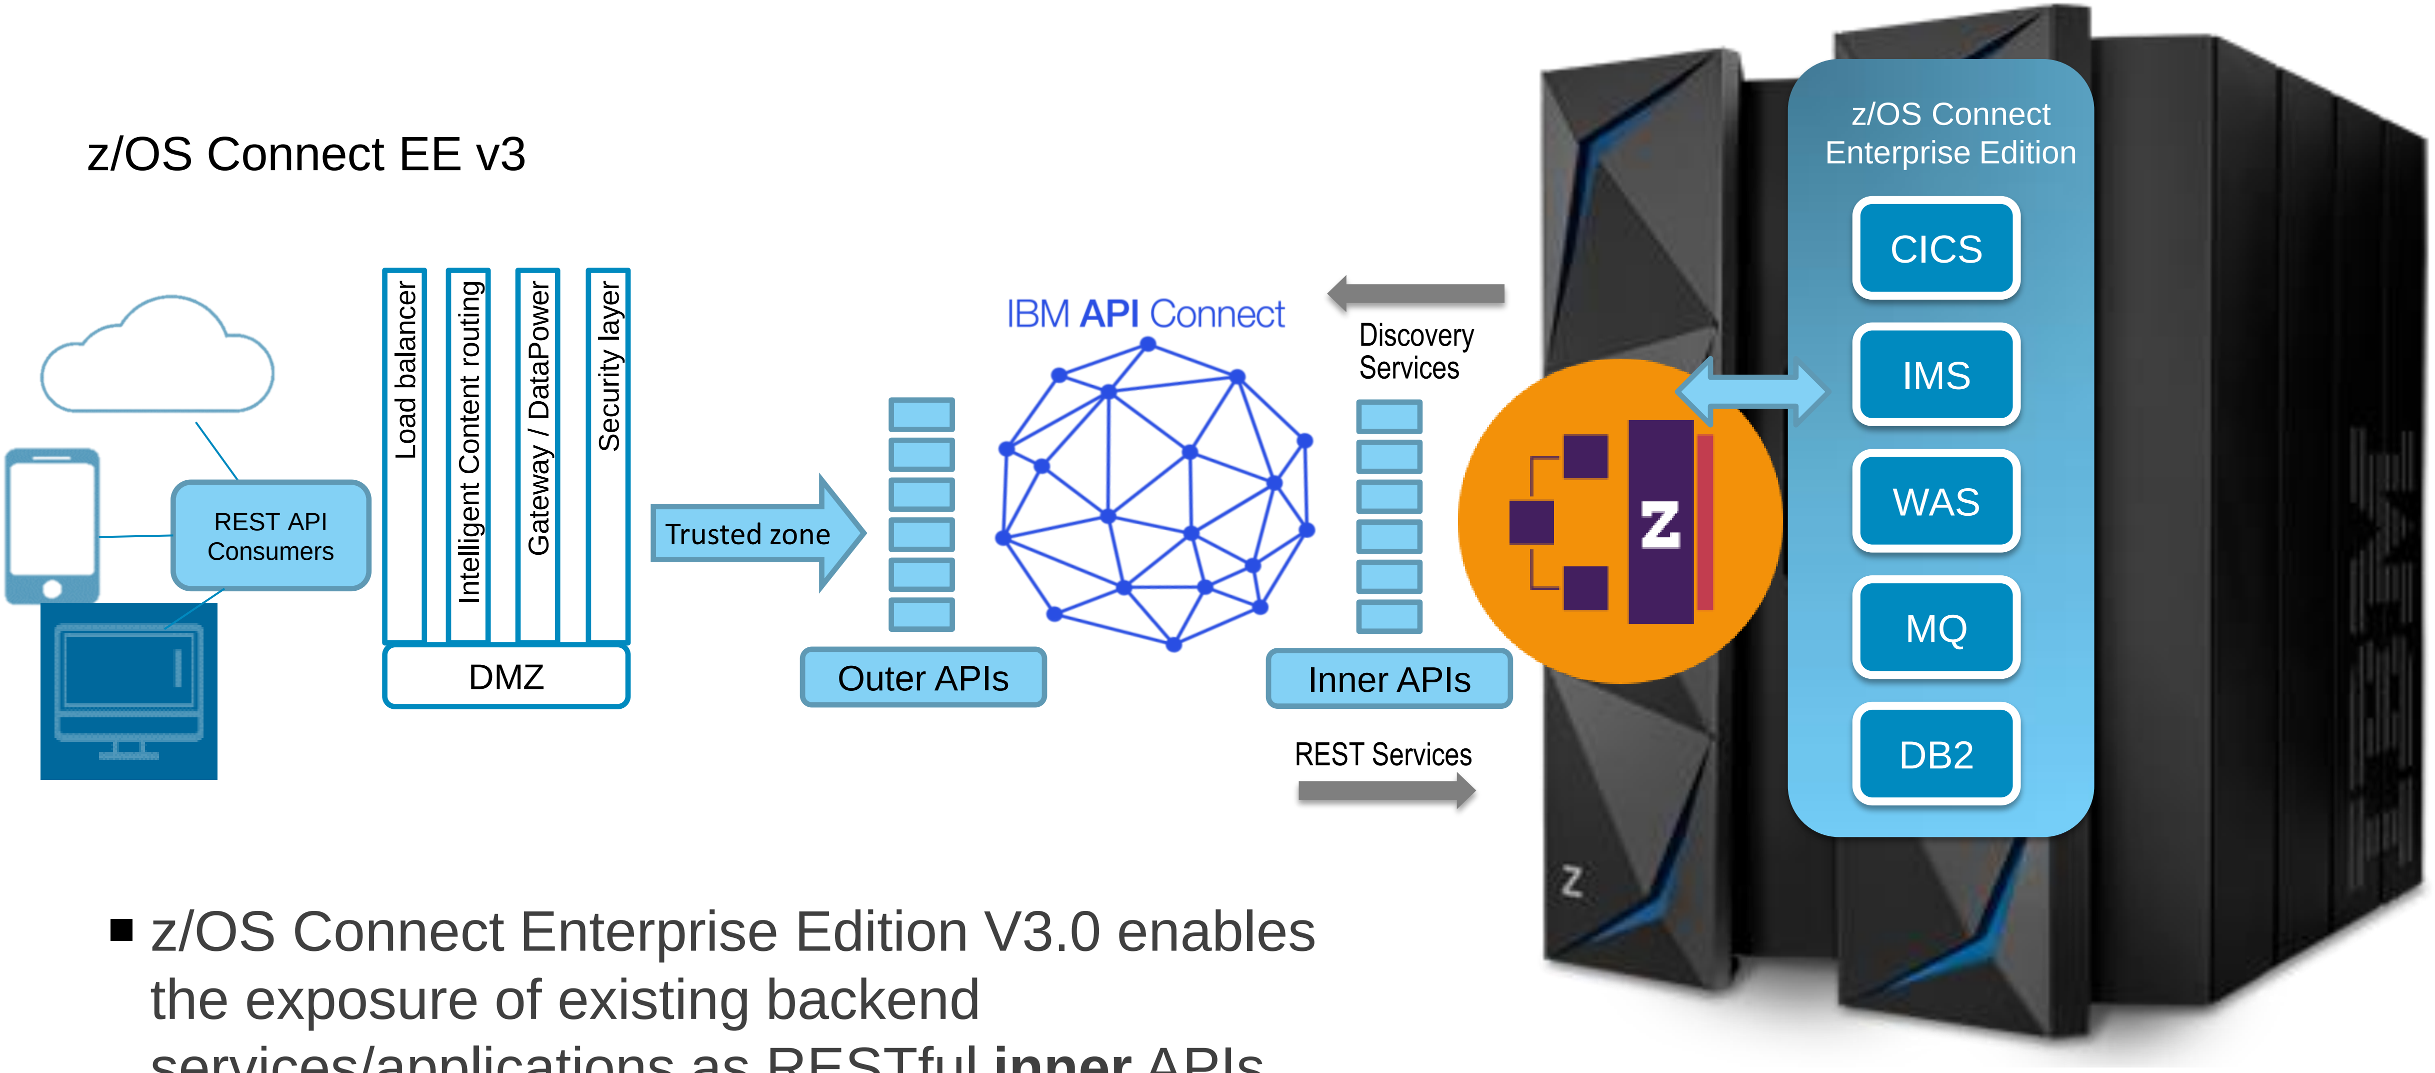

- z/OS Connect Enterprise Edition V3.0 enables the exposure of existing backend services/applications as RESTful **inner** APIs

Data Virtualization for Mainframe API Enablement

z/OS Connect EE v3 and Data Virtualization Manager

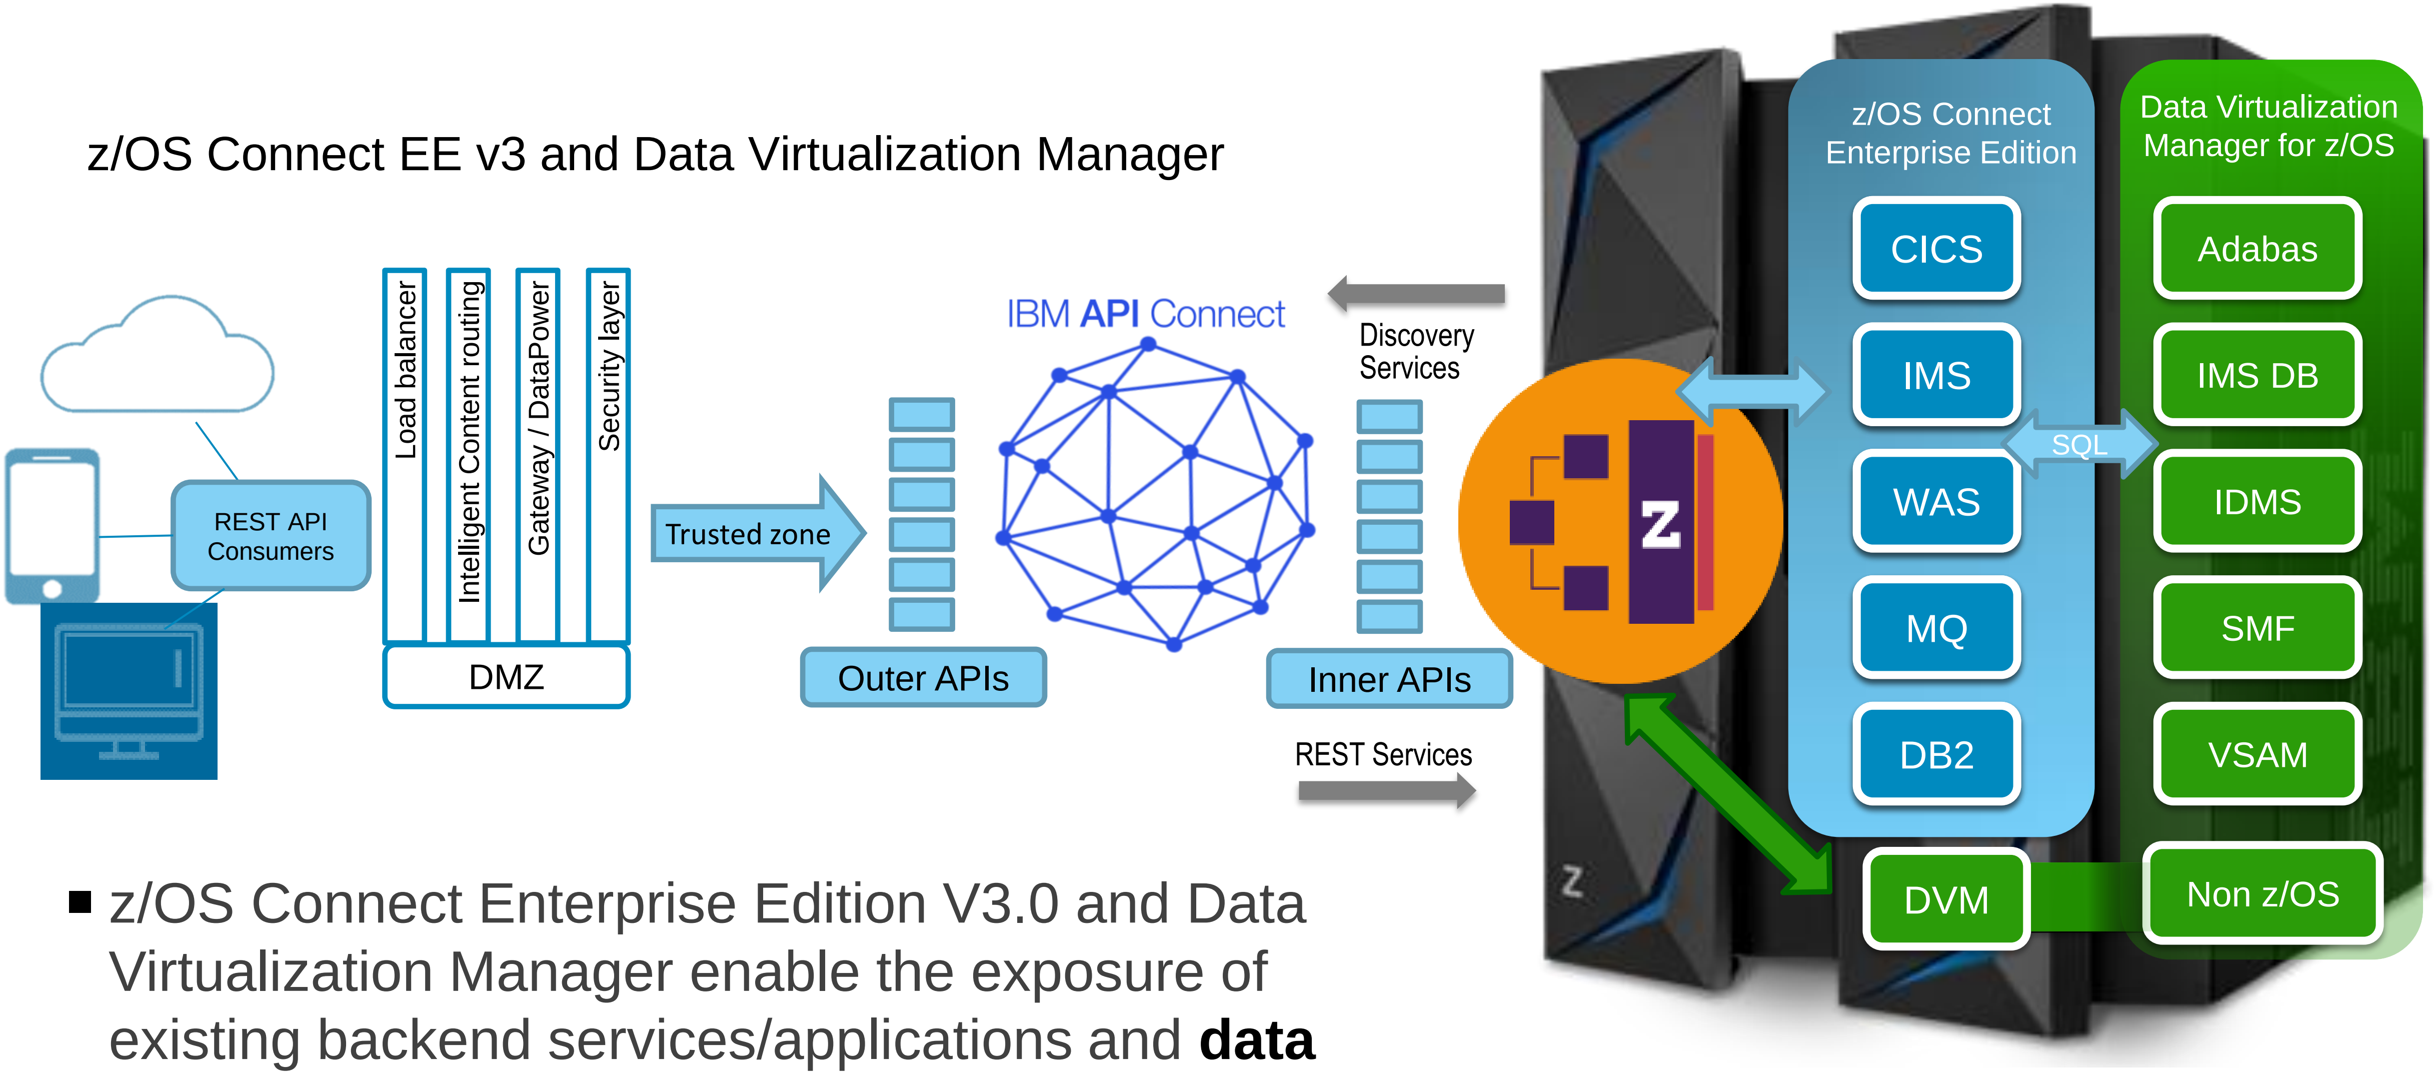

- z/OS Connect Enterprise Edition V3.0 and Data Virtualization Manager enable the exposure of existing backend services/applications and **data** as RESTful inner APIs

IBM Data Virtualization Manager for z/OS

Installable components

- DVM Server (SMP/E)
- Studio (Eclipse-based)
- Drivers (ODBC/JDBC)

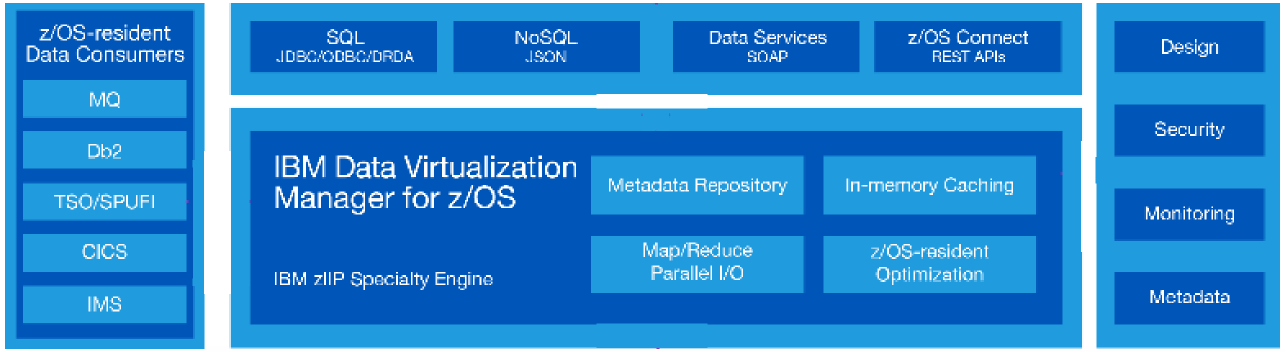

z/OS-resident Data Consumers
MQ
Db2
TSO/SPUFI
CICS
IMS

SQL JDBC/ODBC/DRDA
NoSQL JSON
Data Services SOAP
z/OS Connect REST APIs

IBM Data Virtualization Manager for z/OS
IBM zIIP Specialty Engine
Metadata Repository
In-memory Caching
Map/Reduce Parallel I/O
z/OS-resident Optimization

Design
Security
Monitoring
Metadata

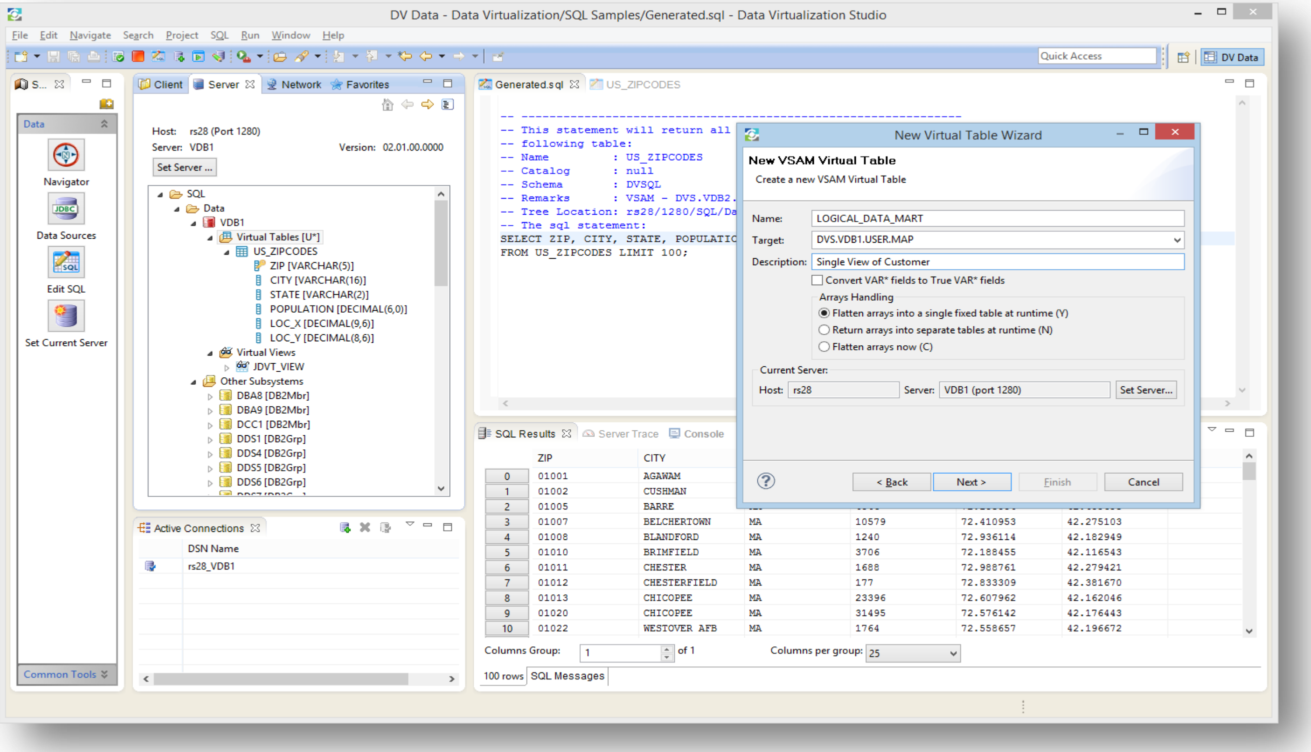

Host: rs28 (Port 1280)
Server: VDB1 Version: 02.01.00.0000
Set Server ...

- SQL
  - Data
    - VDB1
      - Virtual Tables [U*]
        - US_ZIPCODES
          - ZIP [VARCHAR(5)]
          - CITY [VARCHAR(16)]
          - STATE [VARCHAR(2)]
          - POPULATION [DECIMAL(6,0)]
          - LOC_X [DECIMAL(9,6)]
          - LOC_Y [DECIMAL(8,6)]
      - Virtual Views
        - JDVT_VIEW
    - Other Subsystems
      - DBA8 [DB2Mbr]
      - DBA9 [DB2Mbr]
      - DCC1 [DB2Mbr]
      - DDS1 [DB2Grp]
      - DDS4 [DB2Grp]
      - DDS5 [DB2Grp]
      - DDS6 [DB2Grp]

DSN Name
rs28_VDB1

```
-- ----------------------------------------------------------------
-- This statement will return all
-- following table:
-- Name         : US_ZIPCODES
-- Catalog      : null
-- Schema       : DVSQL
-- Remarks      : VSAM - DVS.VDB2.
-- Tree Location: rs28/1280/SQL/Da
-- The sql statement:
SELECT ZIP, CITY, STATE, POPULATIO
FROM US_ZIPCODES LIMIT 100;
```

New Virtual Table Wizard

New VSAM Virtual Table
Create a new VSAM Virtual Table

Name: LOGICAL_DATA_MART
Target: DVS.VDB1.USER.MAP
Description: Single View of Customer
□ Convert VAR* fields to True VAR* fields
Arrays Handling
● Flatten arrays into a single fixed table at runtime (Y)
○ Return arrays into separate tables at runtime (N)
○ Flatten arrays now (C)
Current Server:
Host: rs28 Server: VDB1 (port 1280) Set Server...

< Back | Next > | Finish | Cancel

	ZIP	CITY				
0	01001	AGAWAM				
1	01002	CUSHMAN				
2	01005	BARRE				
3	01007	BELCHERTOWN	MA	10579	72.410953	42.275103
4	01008	BLANDFORD	MA	1240	72.936114	42.182949
5	01010	BRIMFIELD	MA	3706	72.188455	42.116543
6	01011	CHESTER	MA	1688	72.988761	42.279421
7	01012	CHESTERFIELD	MA	177	72.833309	42.381670
8	01013	CHICOPEE	MA	23396	72.607962	42.162046
9	01020	CHICOPEE	MA	31495	72.576142	42.176443
10	01022	WESTOVER AFB	MA	1764	72.558657	42.196672

Columns Group: 1 of 1 Columns per group: 25
100 rows SQL Messages

Installing the IBM Data Virtualization Manager for z/OS (DVM)

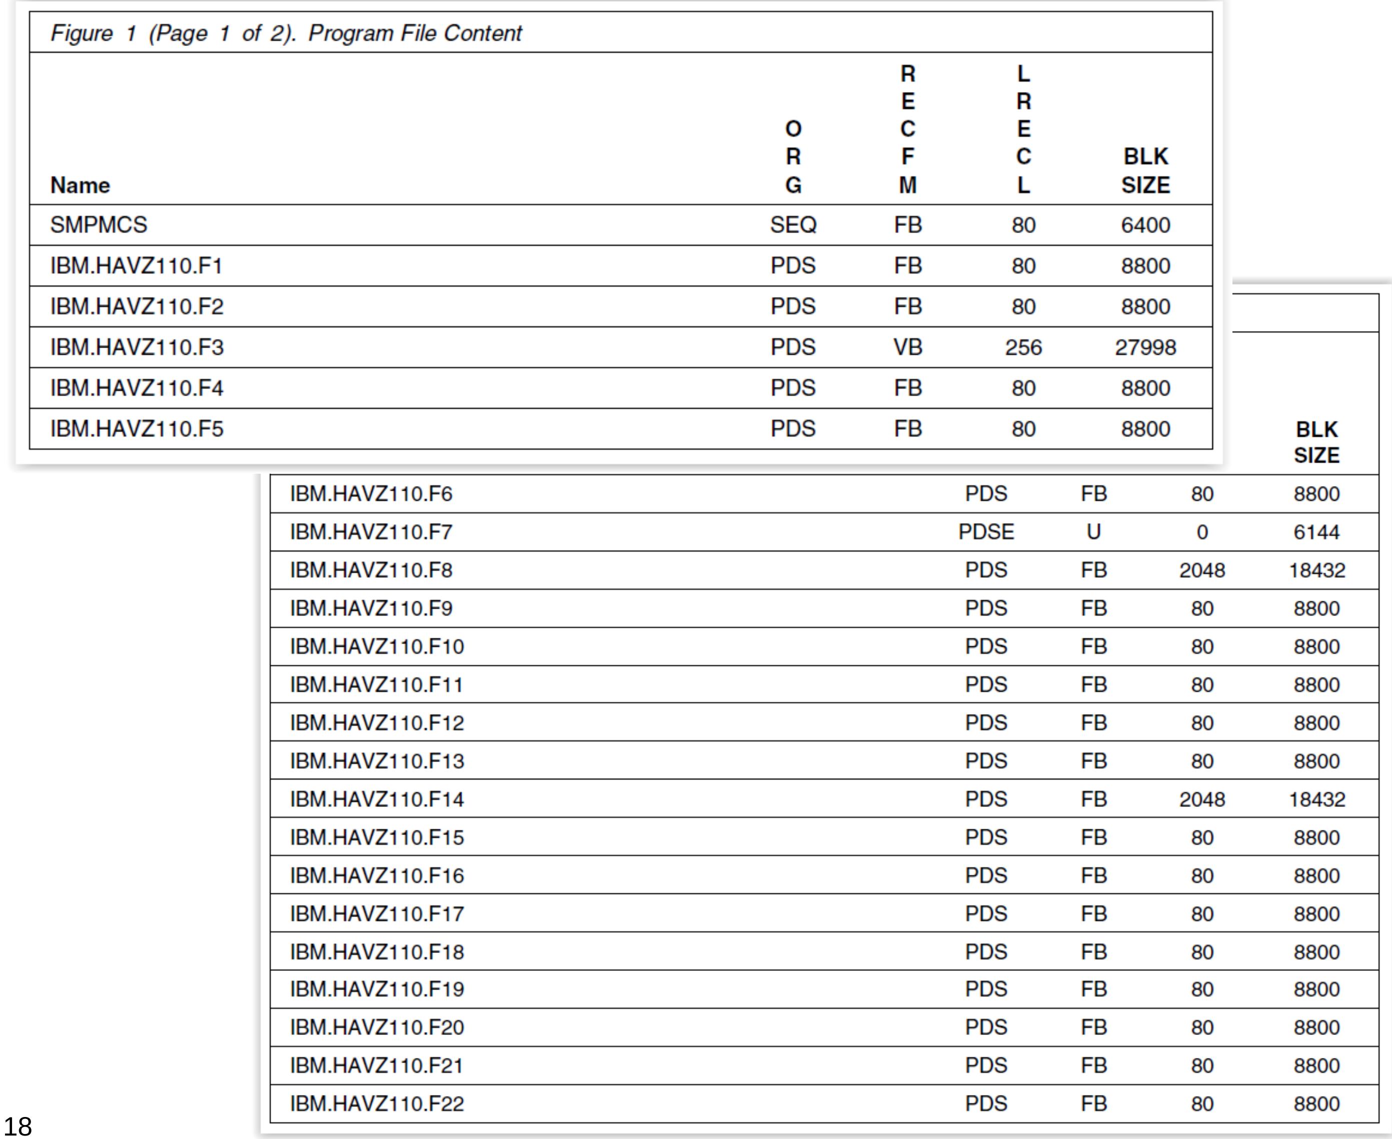

Figure 1 (Page 1 of 2). Program File Content

Name	ORG	RECFM	LRECL	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.HAVZ110.F1	PDS	FB	80	8800
IBM.HAVZ110.F2	PDS	FB	80	8800
IBM.HAVZ110.F3	PDS	VB	256	27998
IBM.HAVZ110.F4	PDS	FB	80	8800
IBM.HAVZ110.F5	PDS	FB	80	8800
IBM.HAVZ110.F6	PDS	FB	80	8800
IBM.HAVZ110.F7	PDSE	U	0	6144
IBM.HAVZ110.F8	PDS	FB	2048	18432
IBM.HAVZ110.F9	PDS	FB	80	8800
IBM.HAVZ110.F10	PDS	FB	80	8800
IBM.HAVZ110.F11	PDS	FB	80	8800
IBM.HAVZ110.F12	PDS	FB	80	8800
IBM.HAVZ110.F13	PDS	FB	80	8800
IBM.HAVZ110.F14	PDS	FB	2048	18432
IBM.HAVZ110.F15	PDS	FB	80	8800
IBM.HAVZ110.F16	PDS	FB	80	8800
IBM.HAVZ110.F17	PDS	FB	80	8800
IBM.HAVZ110.F18	PDS	FB	80	8800
IBM.HAVZ110.F19	PDS	FB	80	8800
IBM.HAVZ110.F20	PDS	FB	80	8800
IBM.HAVZ110.F21	PDS	FB	80	8800
IBM.HAVZ110.F22	PDS	FB	80	8800

Standard SMP/E Installation

Program Directory lists the 22 Rel files shipped with the product

SAVZLOAD, SAVZCLOD and SAVZPRC must be PDSEs

Configuring the DVM Server

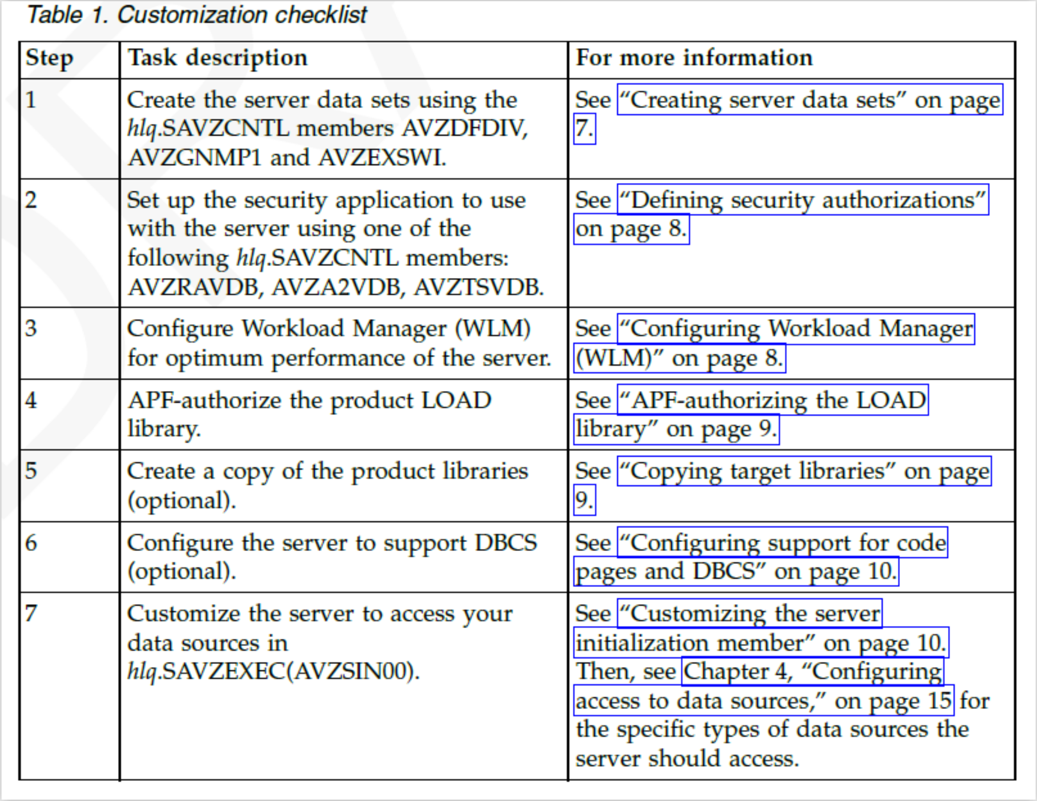

Table 1. Customization checklist

Step	Task description	For more information
1	Create the server data sets using the *hlq*.SAVZCNTL members AVZDFDIV, AVZGNMP1 and AVZEXSWI.	See "Creating server data sets" on page 7.
2	Set up the security application to use with the server using one of the following *hlq*.SAVZCNTL members: AVZRAVDB, AVZA2VDB, AVZTSVDB.	See "Defining security authorizations" on page 8.
3	Configure Workload Manager (WLM) for optimum performance of the server.	See "Configuring Workload Manager (WLM)" on page 8.
4	APF-authorize the product LOAD library.	See "APF-authorizing the LOAD library" on page 9.
5	Create a copy of the product libraries (optional).	See "Copying target libraries" on page 9.
6	Configure the server to support DBCS (optional).	See "Configuring support for code pages and DBCS" on page 10.
7	Customize the server to access your data sources in *hlq*.SAVZEXEC(AVZSIN00).	See "Customizing the server initialization member" on page 10. Then, see Chapter 4, "Configuring access to data sources," on page 15 for the specific types of data sources the server should access.

Installation and Customization Guide contains high level checklist

7 Required tasks

3 Optional tasks

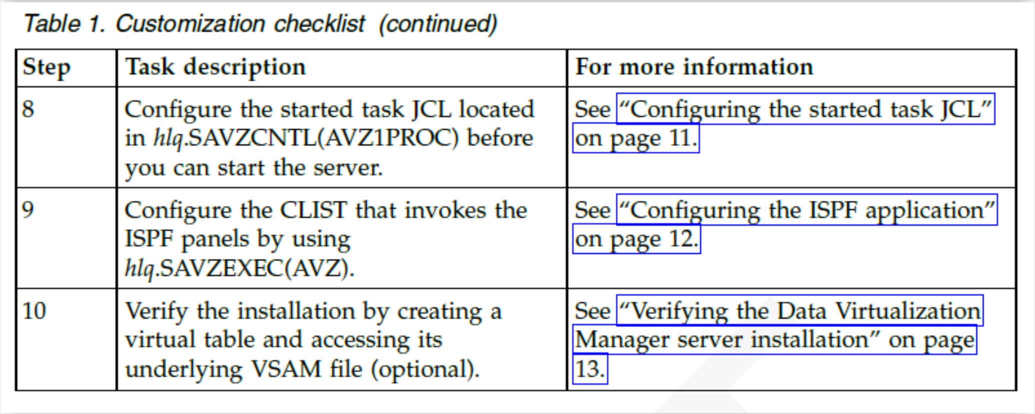

Table 1. Customization checklist (continued)

Step	Task description	For more information
8	Configure the started task JCL located in *hlq*.SAVZCNTL(AVZ1PROC) before you can start the server.	See "Configuring the started task JCL" on page 11.
9	Configure the CLIST that invokes the ISPF panels by using *hlq*.SAVZEXEC(AVZ).	See "Configuring the ISPF application" on page 12.
10	Verify the installation by creating a virtual table and accessing its underlying VSAM file (optional).	See "Verifying the Data Virtualization Manager server installation" on page 13.

DVM Studio connections

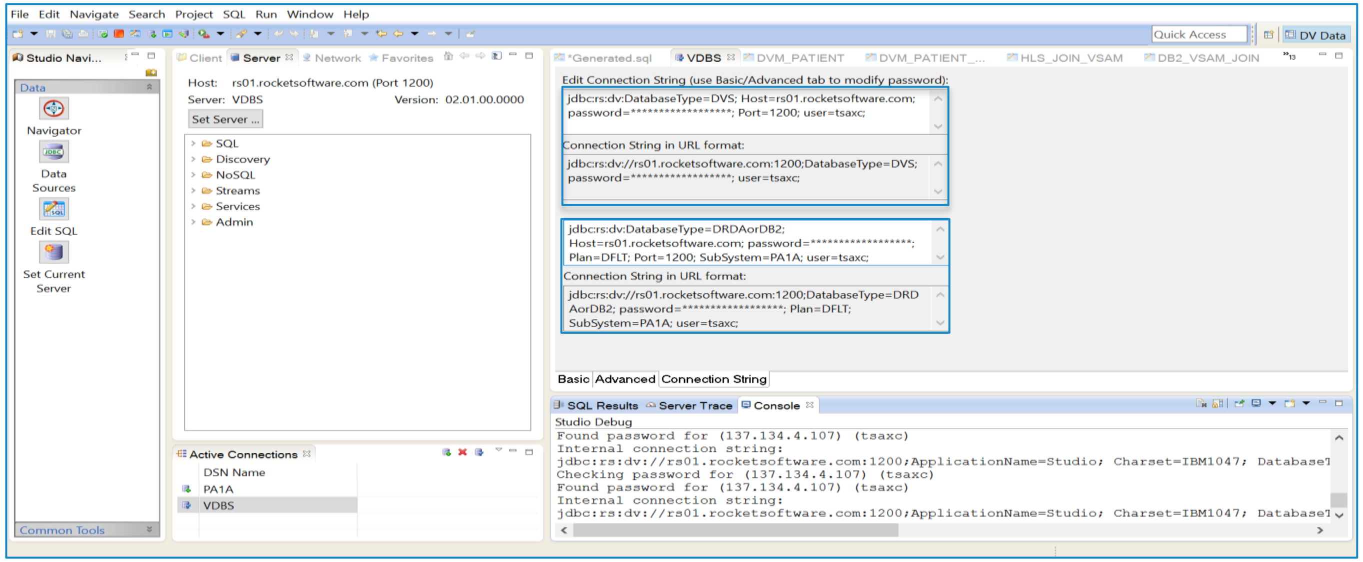

Host: rs01.rocketsoftware.com (Port 1200)

Server: VDBS Version: 02.01.00.0000

DSN Name
PA1A
VDBS

Edit Connection String (use Basic/Advanced tab to modify password):

jdbc:rs:dv:DatabaseType=DVS; Host=rs01.rocketsoftware.com; password=******************; Port=1200; user=tsaxc;

Connection String in URL format:

jdbc:rs:dv://rs01.rocketsoftware.com:1200;DatabaseType=DVS; password=******************; user=tsaxc;

jdbc:rs:dv:DatabaseType=DRDAorDB2; Host=rs01.rocketsoftware.com; password=******************; Plan=DFLT; Port=1200; SubSystem=PA1A; user=tsaxc;

Connection String in URL format:

jdbc:rs:dv://rs01.rocketsoftware.com:1200;DatabaseType=DRDAorDB2; password=******************; Plan=DFLT; SubSystem=PA1A; user=tsaxc;

Studio Debug

```
Found password for (137.134.4.107) (tsaxc)
Internal connection string:
jdbc:rs:dv://rs01.rocketsoftware.com:1200;ApplicationName=Studio; Charset=IBM1047; DatabaseT
Checking password for (137.134.4.107) (tsaxc)
Found password for (137.134.4.107) (tsaxc)
Internal connection string:
jdbc:rs:dv://rs01.rocketsoftware.com:1200;ApplicationName=Studio; Charset=IBM1047; DatabaseT
```

- Connection String
  - JDBC Connection string to DVM Started Task on z/OS
  - JDBC Connection string to DB2

```
  SDSF OUTPUT DISPLAY VDBS      STC07224  DSID     2 LINE 110
  COMMAND INPUT ===>
13.33.32 STC07224  VDS1290H VDBS     OEPORTNUMBER           '1200'
```

DVM Studio view

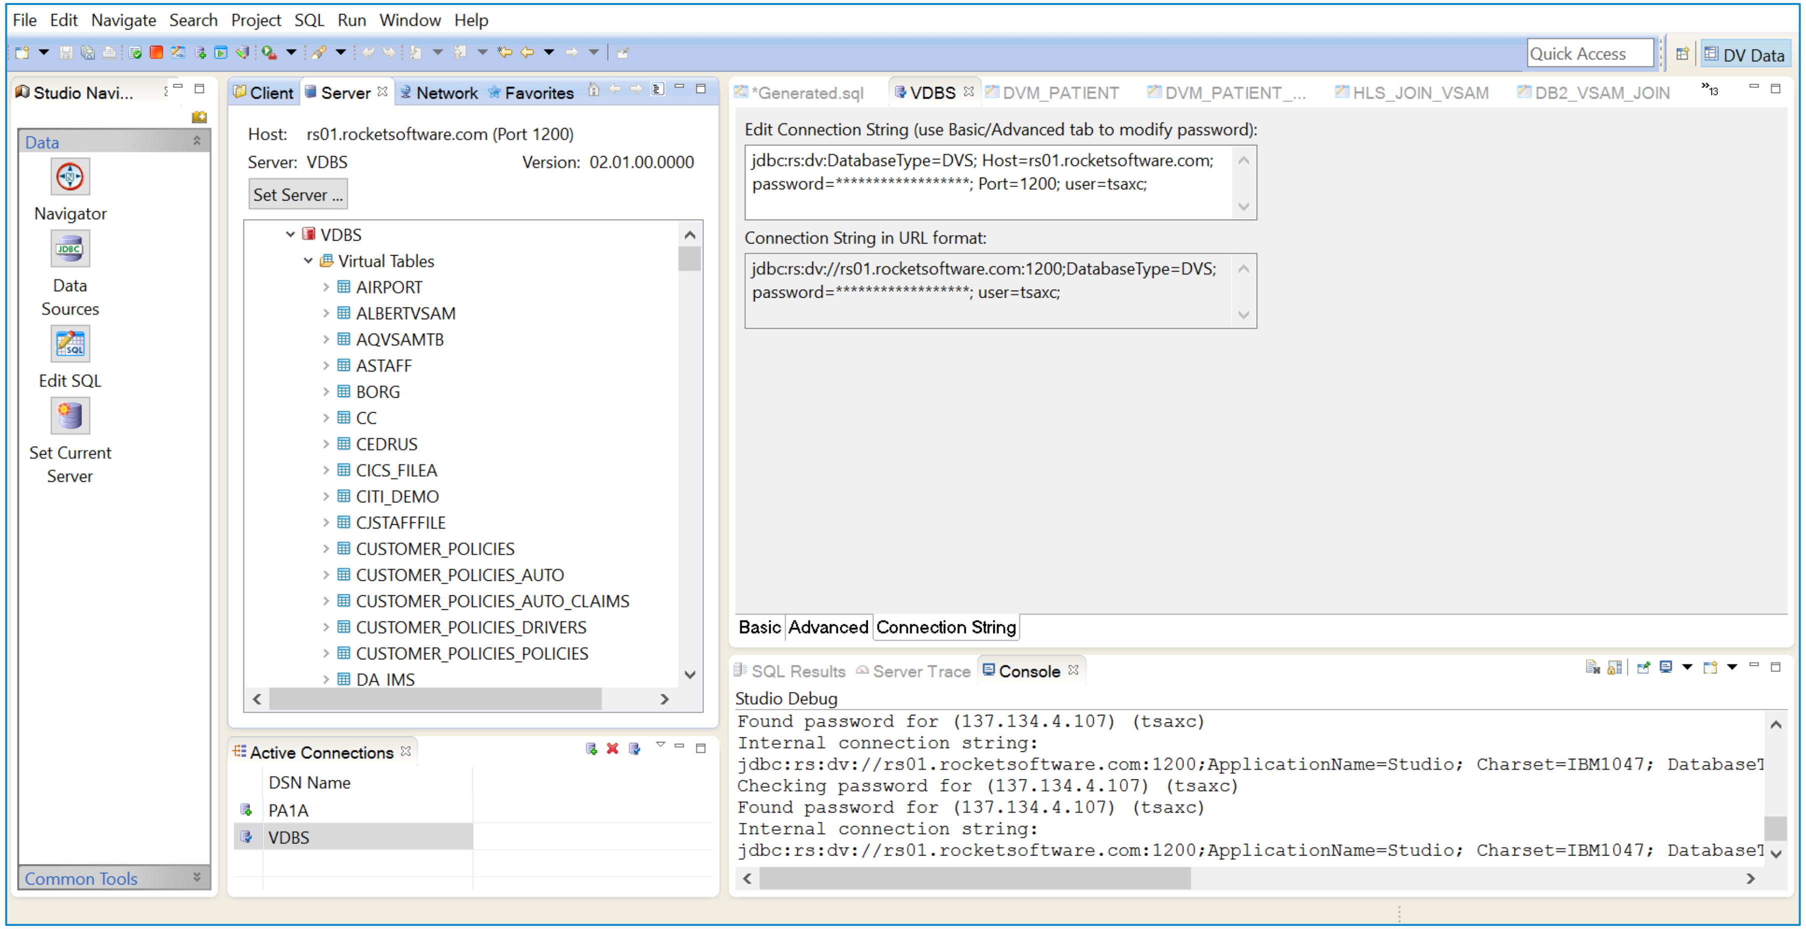

Host: rs01.rocketsoftware.com (Port 1200)

Server: VDBS Version: 02.01.00.0000

- VDBS
  - Virtual Tables
    - AIRPORT
    - ALBERTVSAM
    - AQVSAMTB
    - ASTAFF
    - BORG
    - CC
    - CEDRUS
    - CICS_FILEA
    - CITI_DEMO
    - CJSTAFFFILE
    - CUSTOMER_POLICIES
    - CUSTOMER_POLICIES_AUTO
    - CUSTOMER_POLICIES_AUTO_CLAIMS
    - CUSTOMER_POLICIES_DRIVERS
    - CUSTOMER_POLICIES_POLICIES
    - DA_IMS

Edit Connection String (use Basic/Advanced tab to modify password):

jdbc:rs:dv:DatabaseType=DVS; Host=rs01.rocketsoftware.com; password=****************; Port=1200; user=tsaxc;

Connection String in URL format:

jdbc:rs:dv://rs01.rocketsoftware.com:1200;DatabaseType=DVS; password=****************; user=tsaxc;

DSN Name
PA1A
VDBS

```
Found password for (137.134.4.107) (tsaxc)
Internal connection string:
jdbc:rs:dv://rs01.rocketsoftware.com:1200;ApplicationName=Studio; Charset=IBM1047; DatabaseT
Checking password for (137.134.4.107) (tsaxc)
Found password for (137.134.4.107) (tsaxc)
Internal connection string:
jdbc:rs:dv://rs01.rocketsoftware.com:1200;ApplicationName=Studio; Charset=IBM1047; DatabaseT
```

- Virtual table view within the studio
  - View existing virtual tables
  - Create new virtual tables
  - Execute sample test queries

Creating virtual tables

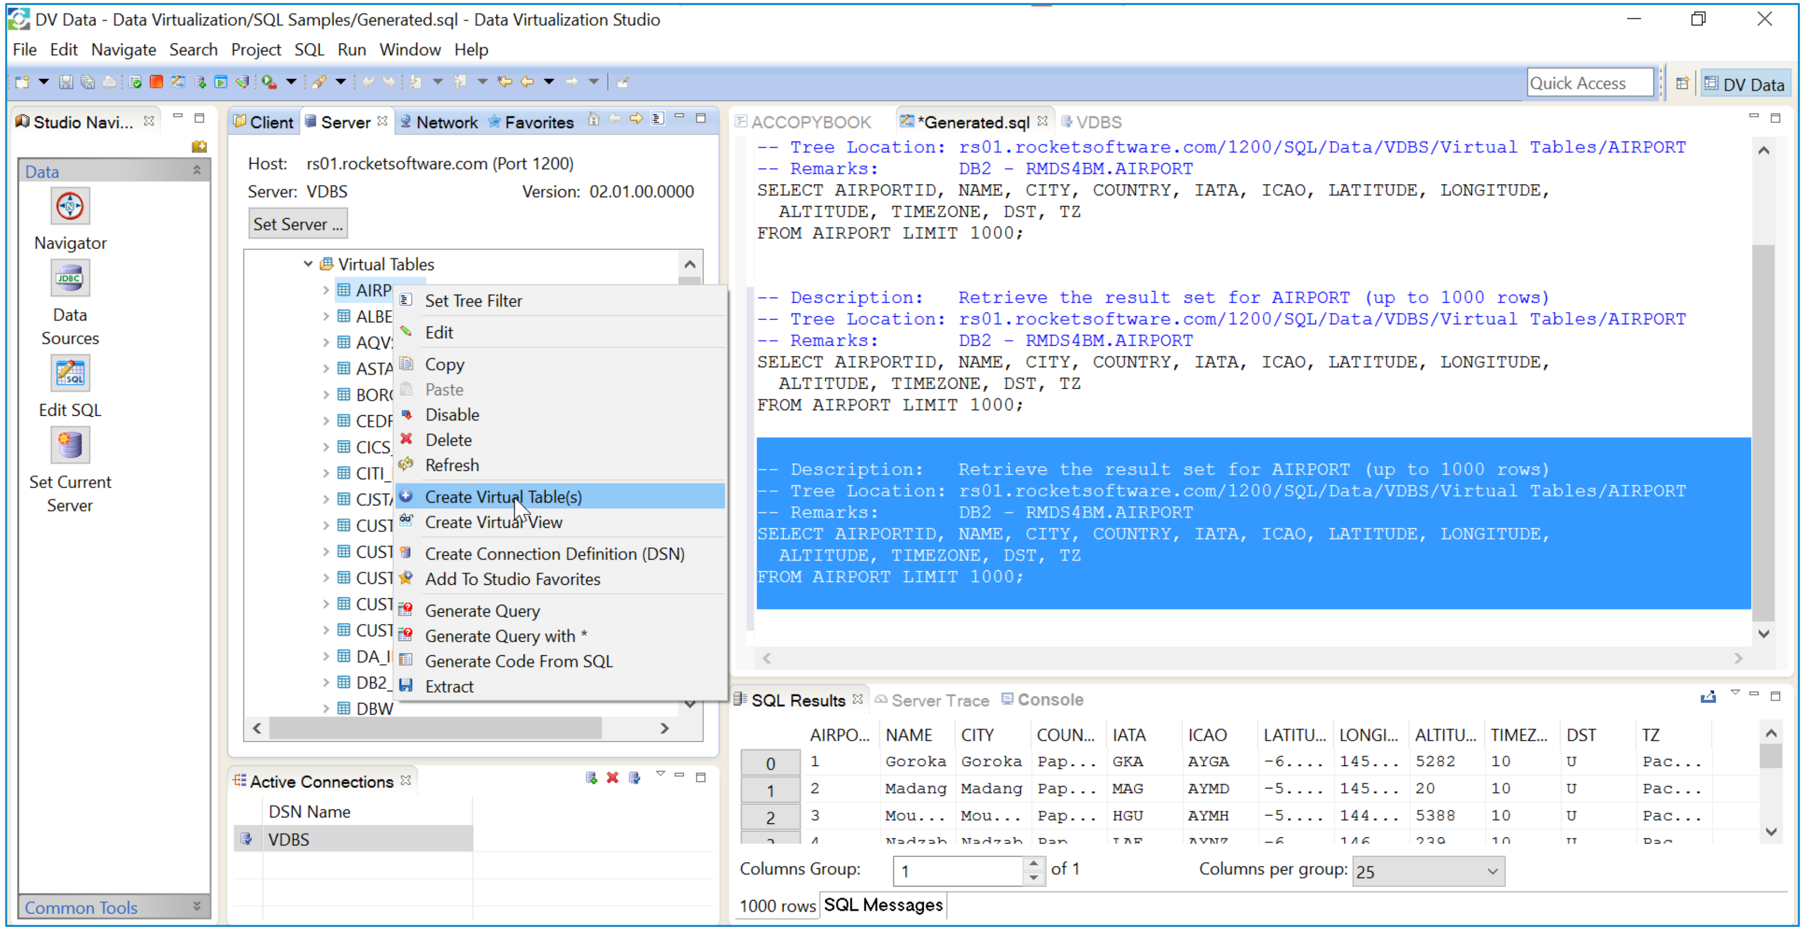

Host: rs01.rocketsoftware.com (Port 1200)

Server: VDBS Version: 02.01.00.0000

```
-- Tree Location: rs01.rocketsoftware.com/1200/SQL/Data/VDBS/Virtual Tables/AIRPORT
-- Remarks:        DB2 - RMDS4BM.AIRPORT
SELECT AIRPORTID, NAME, CITY, COUNTRY, IATA, ICAO, LATITUDE, LONGITUDE,
  ALTITUDE, TIMEZONE, DST, TZ
FROM AIRPORT LIMIT 1000;


-- Description:   Retrieve the result set for AIRPORT (up to 1000 rows)
-- Tree Location: rs01.rocketsoftware.com/1200/SQL/Data/VDBS/Virtual Tables/AIRPORT
-- Remarks:        DB2 - RMDS4BM.AIRPORT
SELECT AIRPORTID, NAME, CITY, COUNTRY, IATA, ICAO, LATITUDE, LONGITUDE,
  ALTITUDE, TIMEZONE, DST, TZ
FROM AIRPORT LIMIT 1000;


-- Description:   Retrieve the result set for AIRPORT (up to 1000 rows)
-- Tree Location: rs01.rocketsoftware.com/1200/SQL/Data/VDBS/Virtual Tables/AIRPORT
-- Remarks:        DB2 - RMDS4BM.AIRPORT
SELECT AIRPORTID, NAME, CITY, COUNTRY, IATA, ICAO, LATITUDE, LONGITUDE,
  ALTITUDE, TIMEZONE, DST, TZ
FROM AIRPORT LIMIT 1000;
```

	AIRPO...	NAME	CITY	COUN...	IATA	ICAO	LATITU...	LONGI...	ALTITU...	TIMEZ...	DST	TZ
0	1	Goroka	Goroka	Pap...	GKA	AYGA	-6....	145...	5282	10	U	Pac...
1	2	Madang	Madang	Pap...	MAG	AYMD	-5....	145...	20	10	U	Pac...
2	3	Mou...	Mou...	Pap...	HGU	AYMH	-5....	144...	5388	10	U	Pac...
3	4	Nadzab	Nadzab	Pap...	LAE	AYNZ	-6....	146...	239	10	U	Pac...

Columns Group: 1 of 1 Columns per group: 25

1000 rows

- Virtual Tables
  - ➢ A right-click supplies various virtual table options
  - ➢ Select the “create virtual tables” option

Data sources

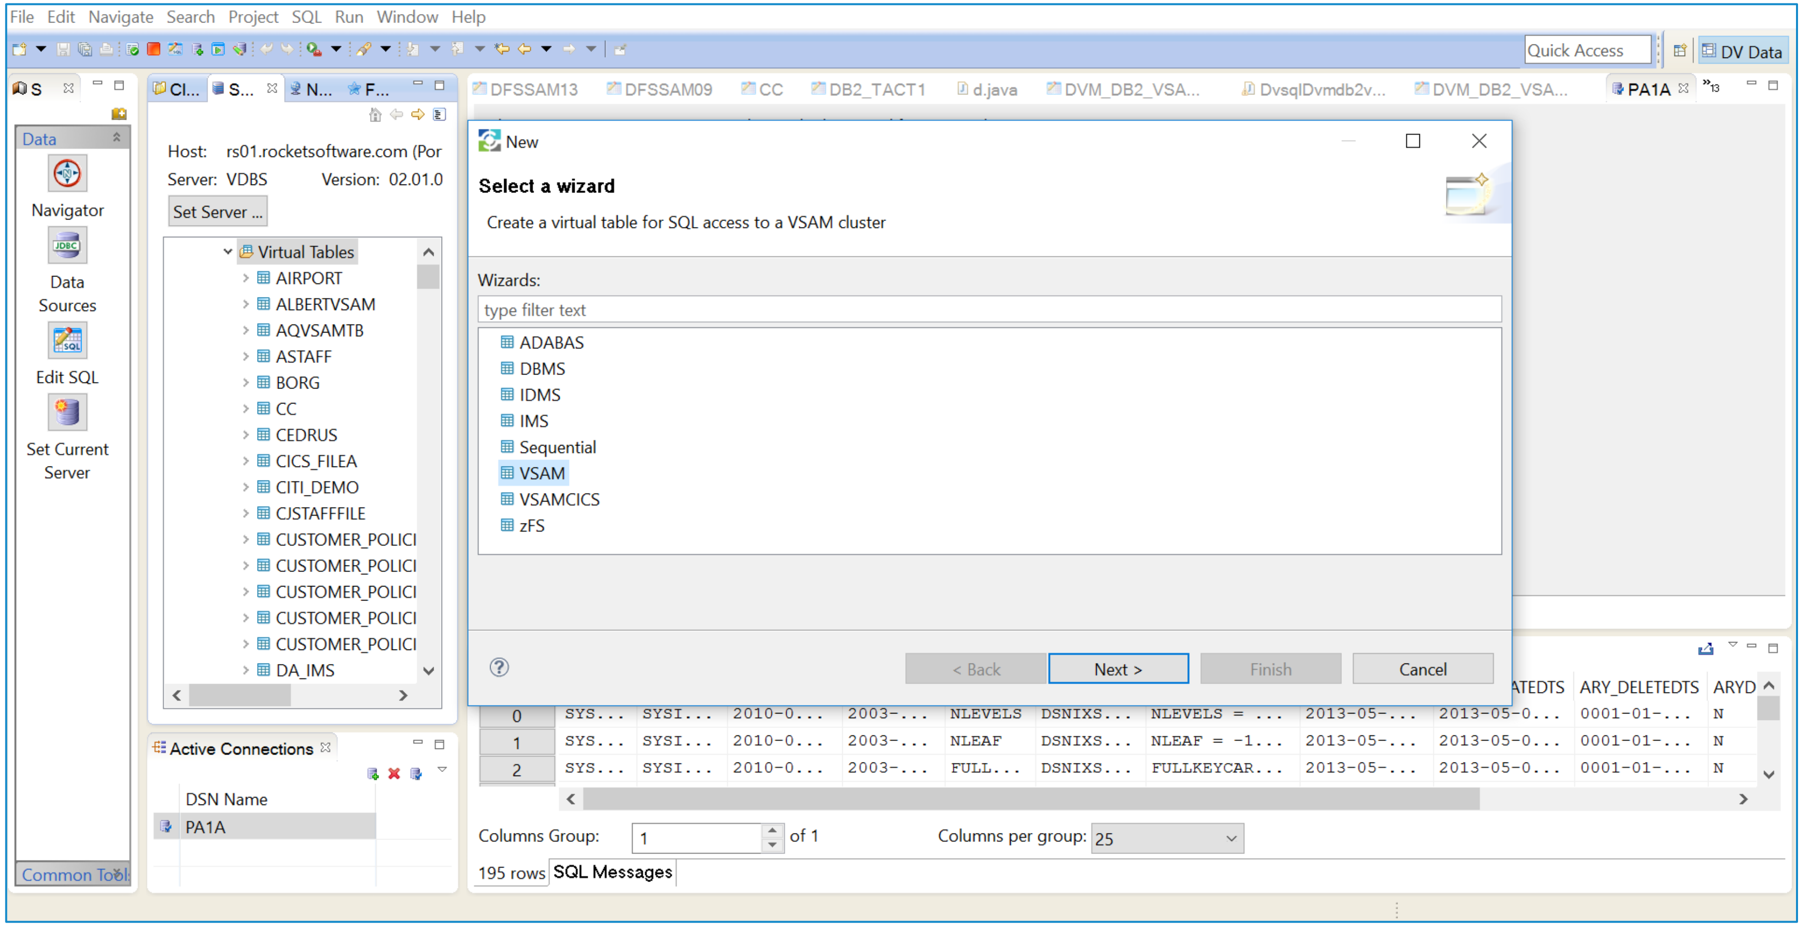

File Edit Navigate Search Project SQL Run Window Help
Quick Access
DV Data
DFSSAM13 DFSSAM09 CC DB2_TACT1 d.java DVM_DB2_VSA... DvsqlDvmdb2v... DVM_DB2_VSA... PA1A
Data
Navigator
Data Sources
Edit SQL
Set Current Server
Common Tool
Cl... S... N... F...

Host: rs01.rocketsoftware.com (Por

Server: VDBS Version: 02.01.0

Set Server ...

- Virtual Tables
  - AIRPORT
  - ALBERTVSAM
  - AQVSAMTB
  - ASTAFF
  - BORG
  - CC
  - CEDRUS
  - CICS_FILEA
  - CITI_DEMO
  - CJSTAFFFILE
  - CUSTOMER_POLICI
  - CUSTOMER_POLICI
  - CUSTOMER_POLICI
  - CUSTOMER_POLICI
  - CUSTOMER_POLICI
  - DA_IMS

Active Connections

DSN Name
PA1A

New

Select a wizard

Create a virtual table for SQL access to a VSAM cluster

Wizards:

type filter text

- ADABAS
- DBMS
- IDMS
- IMS
- Sequential
- VSAM
- VSAMCICS
- zFS

< Back | Next > | Finish | Cancel

									ATEDTS	ARY_DELETEDTS	ARYD
0	SYS...	SYSI...	2010-0...	2003-...	NLEVELS	DSNIXS...	NLEVELS = ...	2013-05-...	2013-05-0...	0001-01-...	N
1	SYS...	SYSI...	2010-0...	2003-...	NLEAF	DSNIXS...	NLEAF = -1...	2013-05-...	2013-05-0...	0001-01-...	N
2	SYS...	SYSI...	2010-0...	2003-...	FULL...	DSNIXS...	FULLKEYCAR...	2013-05-...	2013-05-0...	0001-01-...	N

Columns Group: 1 of 1 Columns per group: 25

195 rows SQL Messages

- Data sources
  - List of available z/OS data sources
  - VSAM will be the data source in this example

Virtual table options

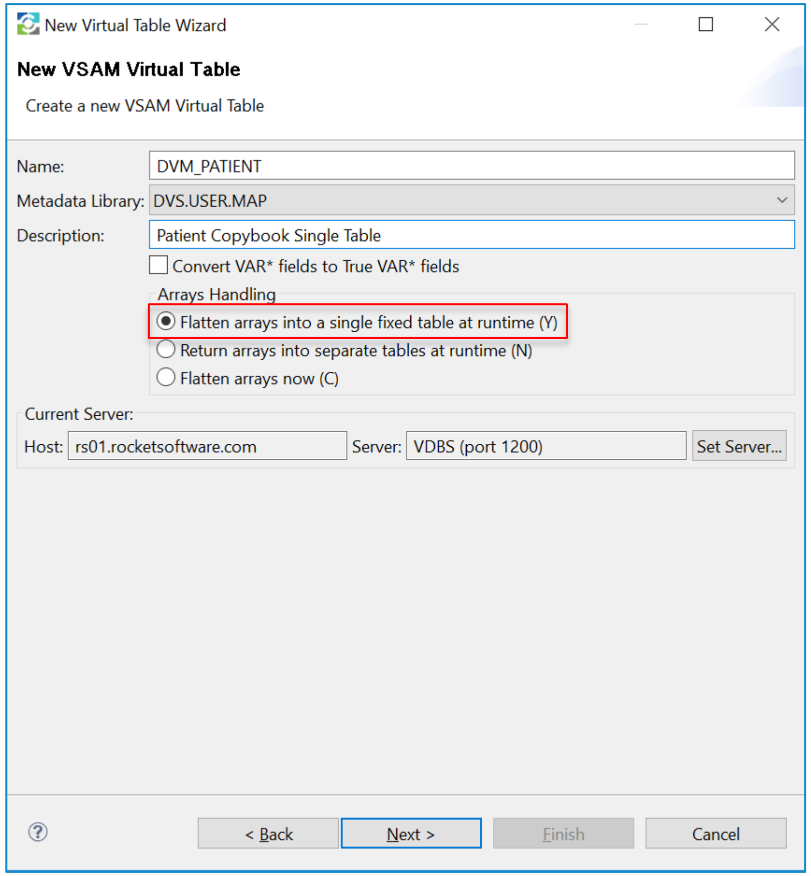

New Virtual Table Wizard

New VSAM Virtual Table

Create a new VSAM Virtual Table

Name: DVM_PATIENT

Metadata Library: DVS.USER.MAP

Description: Patient Copybook Single Table

□ Convert VAR* fields to True VAR* fields

Arrays Handling

● Flatten arrays into a single fixed table at runtime (Y)

○ Return arrays into separate tables at runtime (N)

○ Flatten arrays now (C)

Current Server:

Host: rs01.rocketsoftware.com Server: VDBS (port 1200)

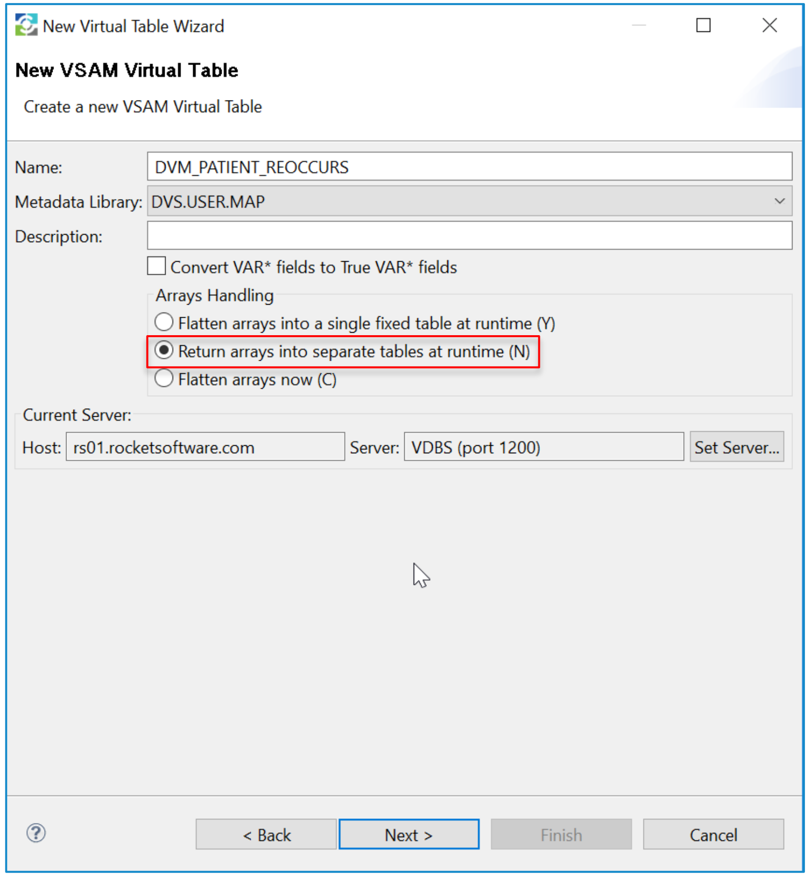

New Virtual Table Wizard

New VSAM Virtual Table

Create a new VSAM Virtual Table

Name: DVM_PATIENT_REOCCURS

Metadata Library: DVS.USER.MAP

Description:

□ Convert VAR* fields to True VAR* fields

Arrays Handling

○ Flatten arrays into a single fixed table at runtime (Y)

● Return arrays into separate tables at runtime (N)

○ Flatten arrays now (C)

Current Server:

Host: rs01.rocketsoftware.com Server: VDBS (port 1200)

- Virtual table options
  - ➢ Supply a name for the virtual tables, names are non-qualified names
  - ➢ Option to create one flat table or parent/child tables for reoccurs
  - ➢ PDSE for metadata library

Copybooks

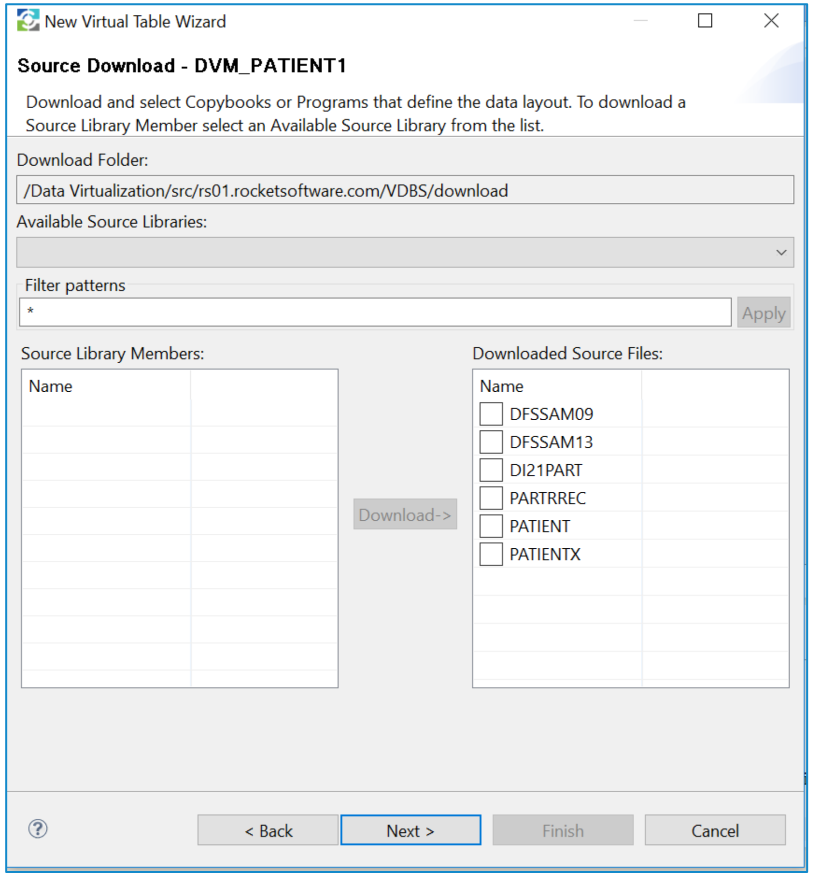

Source Download - DVM_PATIENT1

Download and select Copybooks or Programs that define the data layout. To download a Source Library Member select an Available Source Library from the list.

Download Folder: /Data Virtualization/src/rs01.rocketsoftware.com/VDBS/download

Available Source Libraries:

Filter patterns: *

Source Library Members:

Name

Downloaded Source Files:

Name
□ DFSSAM09
□ DFSSAM13
□ DI21PART
□ PARTRREC
□ PATIENT
□ PATIENTX

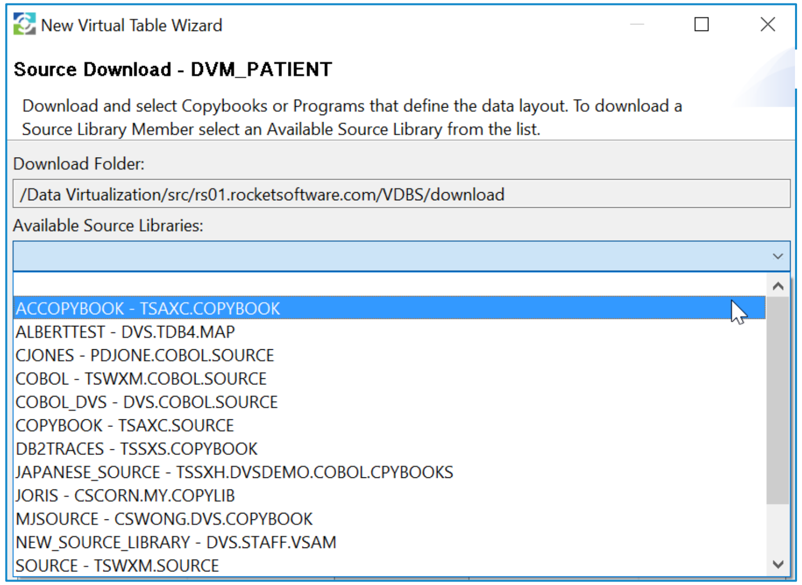

Source Download - DVM_PATIENT

Download and select Copybooks or Programs that define the data layout. To download a Source Library Member select an Available Source Library from the list.

Download Folder: /Data Virtualization/src/rs01.rocketsoftware.com/VDBS/download

Available Source Libraries:

- ACCOPYBOOK - TSAXC.COPYBOOK
- ALBERTTEST - DVS.TDB4.MAP
- CJONES - PDJONE.COBOL.SOURCE
- COBOL - TSWXM.COBOL.SOURCE
- COBOL_DVS - DVS.COBOL.SOURCE
- COPYBOOK - TSAXC.SOURCE
- DB2TRACES - TSSXS.COPYBOOK
- JAPANESE_SOURCE - TSSXH.DVSDEMO.COBOL.CPYBOOKS
- JORIS - CSCORN.MY.COPYLIB
- MJSOURCE - CSWONG.DVS.COPYBOOK
- NEW_SOURCE_LIBRARY - DVS.STAFF.VSAM
- SOURCE - TSWXM.SOURCE

- Copybook options
  - ➢ Pull down menu to choose what copybook dataset to access
  - ➢ Members will be displayed
  - ➢ Select/download the appropriate VSAM copybook

Copybook - continued

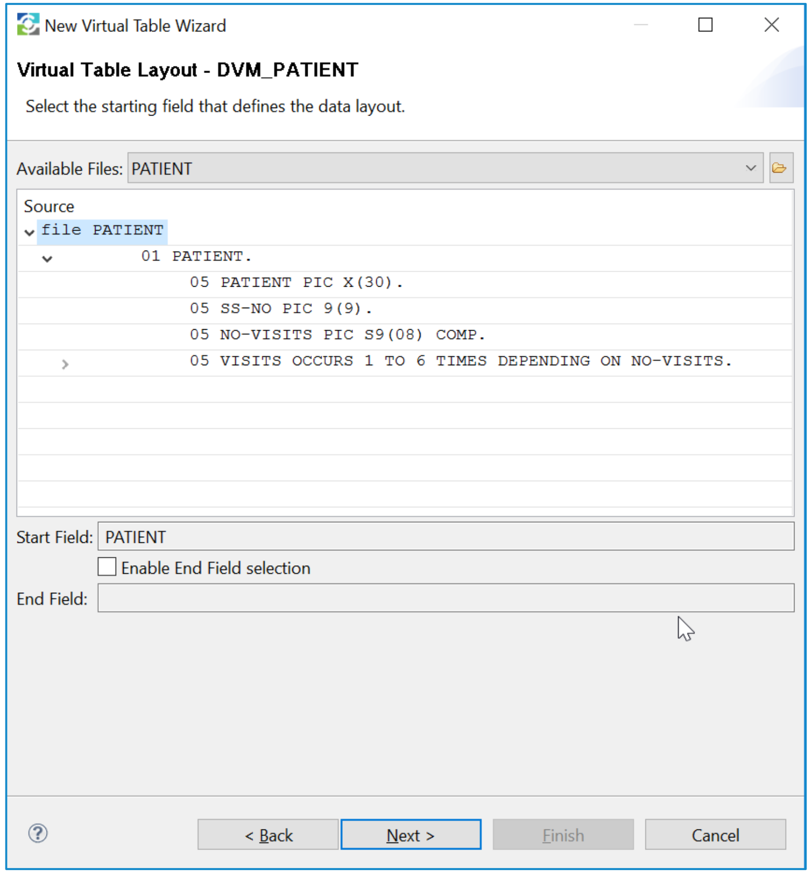

Virtual Table Layout - DVM_PATIENT

Select the starting field that defines the data layout.

Available Files: PATIENT

Source

```
file PATIENT
        01 PATIENT.
            05 PATIENT PIC X(30).
            05 SS-NO PIC 9(9).
            05 NO-VISITS PIC S9(08) COMP.
            05 VISITS OCCURS 1 TO 6 TIMES DEPENDING ON NO-VISITS.
```

Start Field: PATIENT

□ Enable End Field selection

End Field:

- Copybook options
  - Visual representation of selected copybook
  - Ability to select start/end fields
  - Example – not selecting a "Filler" field

VSAM cluster name

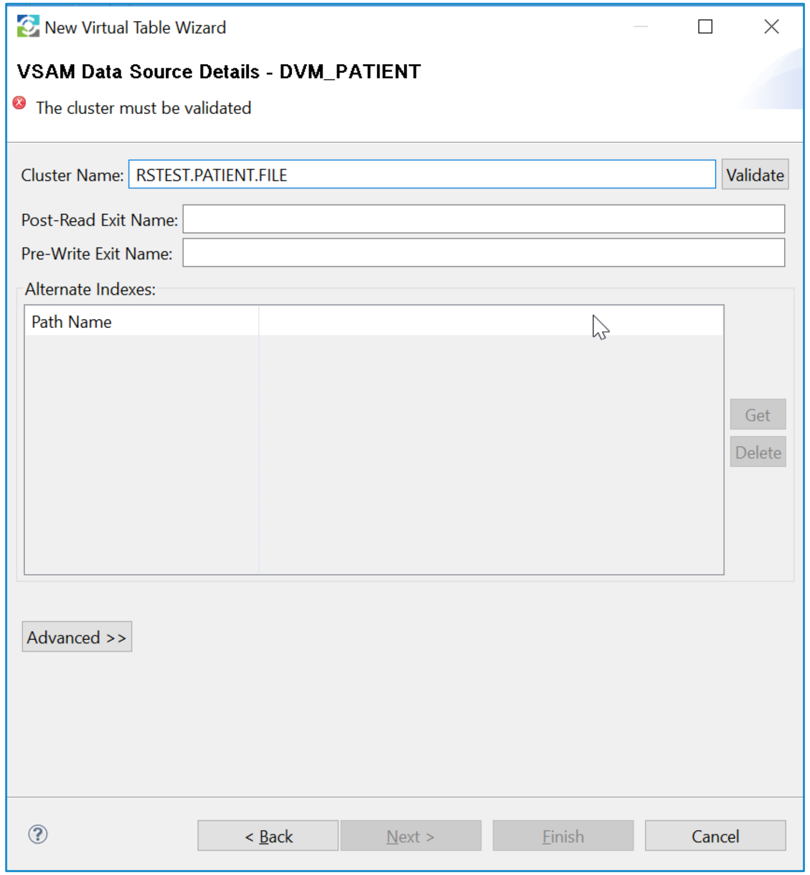

New Virtual Table Wizard

VSAM Data Source Details - DVM_PATIENT

The cluster must be validated

Cluster Name: RSTEST.PATIENT.FILE

Post-Read Exit Name:

Pre-Write Exit Name:

Alternate Indexes:

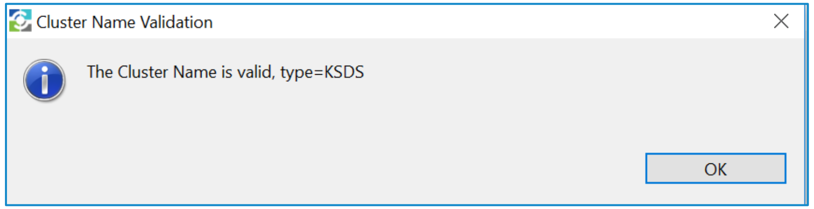

Cluster Name Validation

The Cluster Name is valid, type=KSDS

- Final options
  - ➢ Input VSAM cluster name and click to validate
  - ➢ Post/pre exit options
  - ➢ Ability to use an alternate index
  - ➢ Finish to create

Data validation

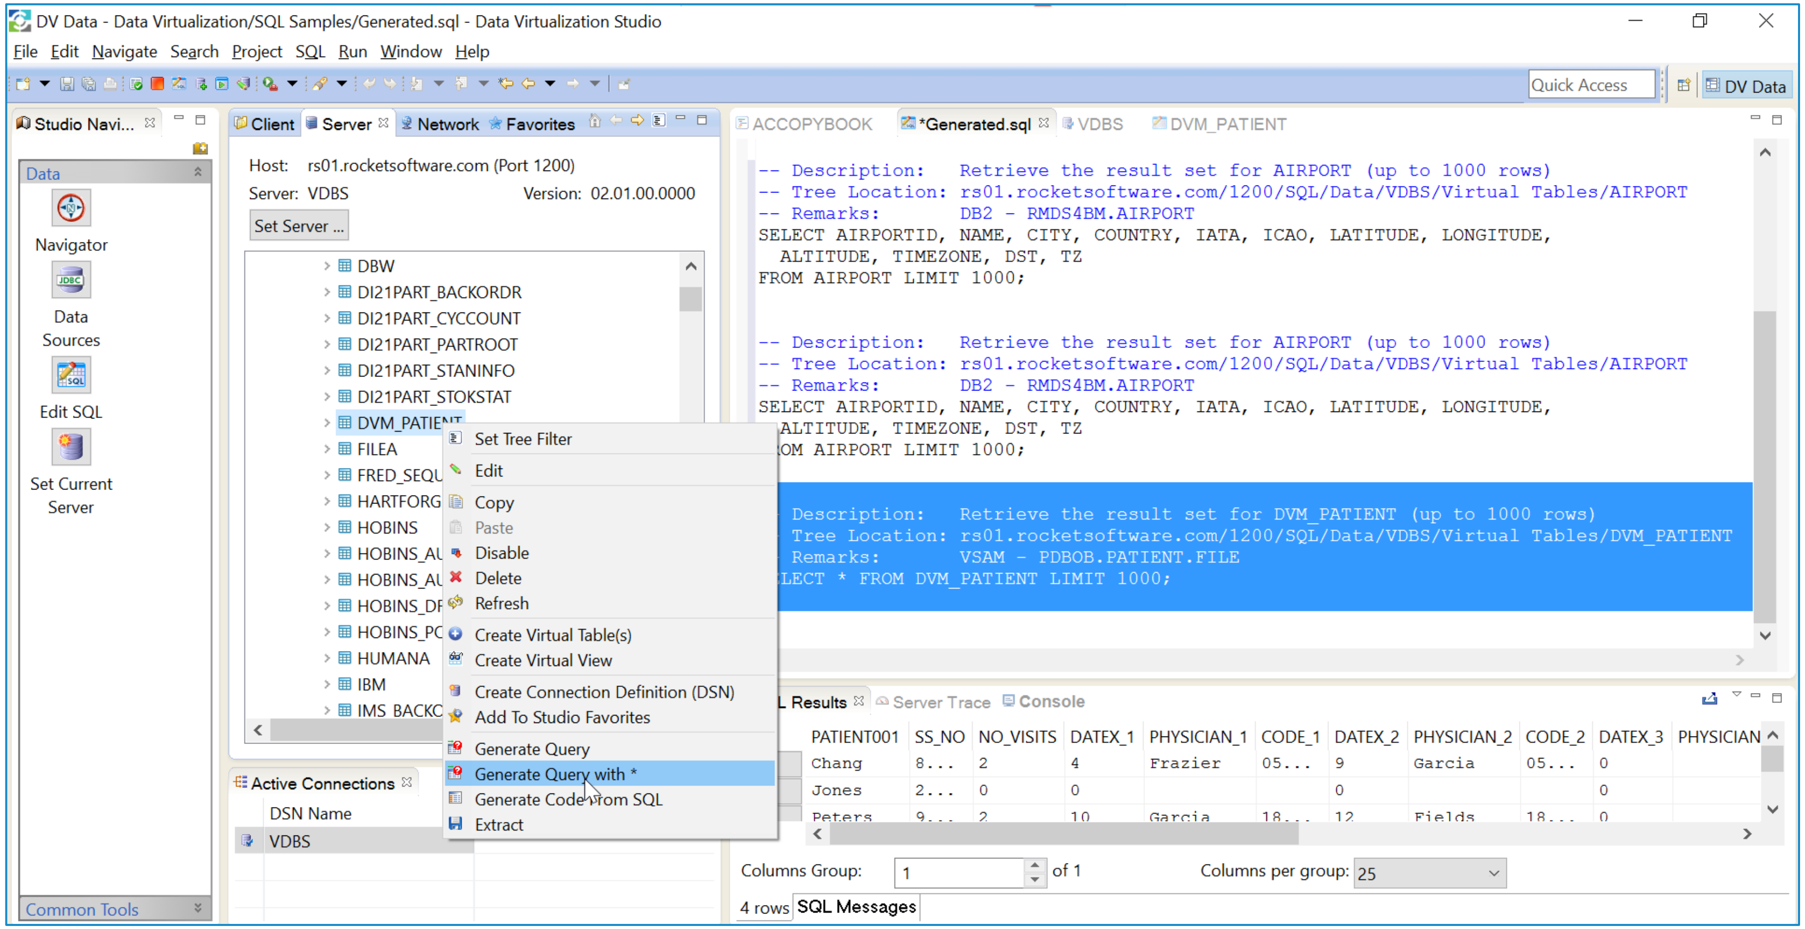

Host: rs01.rocketsoftware.com (Port 1200)
Server: VDBS Version: 02.01.00.0000

```
-- Description:   Retrieve the result set for AIRPORT (up to 1000 rows)
-- Tree Location: rs01.rocketsoftware.com/1200/SQL/Data/VDBS/Virtual Tables/AIRPORT
-- Remarks:       DB2 - RMDS4BM.AIRPORT
SELECT AIRPORTID, NAME, CITY, COUNTRY, IATA, ICAO, LATITUDE, LONGITUDE,
  ALTITUDE, TIMEZONE, DST, TZ
FROM AIRPORT LIMIT 1000;


-- Description:   Retrieve the result set for AIRPORT (up to 1000 rows)
-- Tree Location: rs01.rocketsoftware.com/1200/SQL/Data/VDBS/Virtual Tables/AIRPORT
-- Remarks:       DB2 - RMDS4BM.AIRPORT
SELECT AIRPORTID, NAME, CITY, COUNTRY, IATA, ICAO, LATITUDE, LONGITUDE,
 ALTITUDE, TIMEZONE, DST, TZ
OM AIRPORT LIMIT 1000;


 Description:   Retrieve the result set for DVM_PATIENT (up to 1000 rows)
 Tree Location: rs01.rocketsoftware.com/1200/SQL/Data/VDBS/Virtual Tables/DVM_PATIENT
 Remarks:       VSAM - PDBOB.PATIENT.FILE
LECT * FROM DVM_PATIENT LIMIT 1000;
```

PATIENT001	SS_NO	NO_VISITS	DATEX_1	PHYSICIAN_1	CODE_1	DATEX_2	PHYSICIAN_2	CODE_2	DATEX_3	PHYSICIAN
Chang	8...	2	4	Frazier	05...	9	Garcia	05...	0	
Jones	2...	0	0			0			0	
Peters	9...	2	10	Garcia	18...	12	Fields	18...	0	

Columns Group: 1 of 1 Columns per group: 25

4 rows

- Data validation
  - ➢ Right-click the virtual table name
  - ➢ Select the option to query the data
  - ➢ A SQL virtual table query will be generated to validate the data at rest

Data validation - continued

1

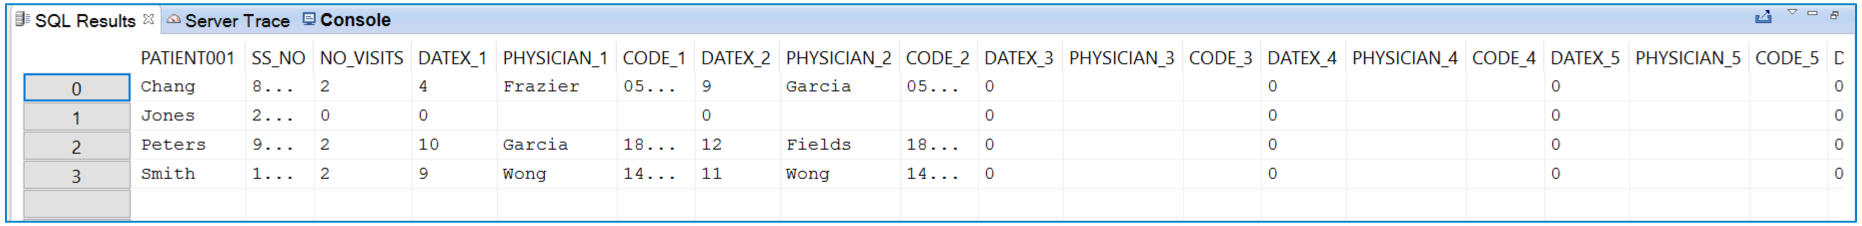

	PATIENT001	SS_NO	NO_VISITS	DATEX_1	PHYSICIAN_1	CODE_1	DATEX_2	PHYSICIAN_2	CODE_2	DATEX_3	PHYSICIAN_3	CODE_3	DATEX_4	PHYSICIAN_4	CODE_4	DATEX_5	PHYSICIAN_5	CODE_5	D
0	Chang	8...	2	4	Frazier	05...	9	Garcia	05...	0			0			0			0
1	Jones	2...	0	0			0			0			0			0			0
2	Peters	9...	2	10	Garcia	18...	12	Fields	18...	0			0			0			0
3	Smith	1...	2	9	Wong	14...	11	Wong	14...	0			0			0			0

2

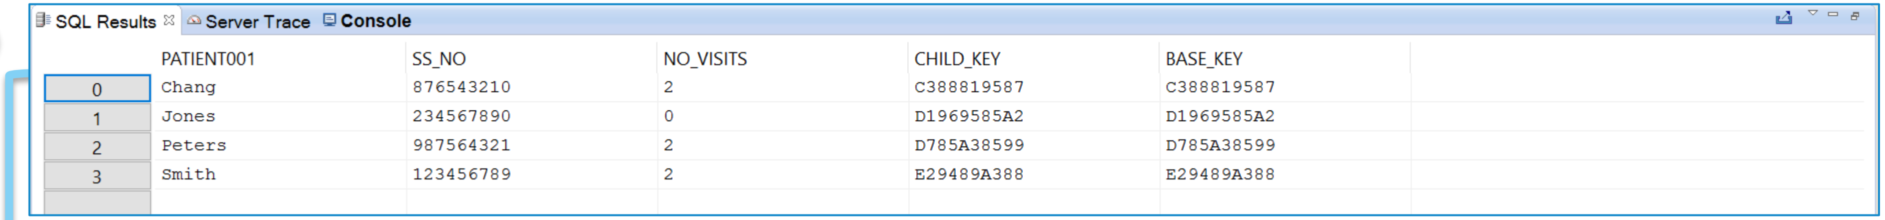

	PATIENT001	SS_NO	NO_VISITS	CHILD_KEY	BASE_KEY
0	Chang	876543210	2	C388819587	C388819587
1	Jones	234567890	0	D1969585A2	D1969585A2
2	Peters	987564321	2	D785A38599	D785A38599
3	Smith	123456789	2	E29489A388	E29489A388

3

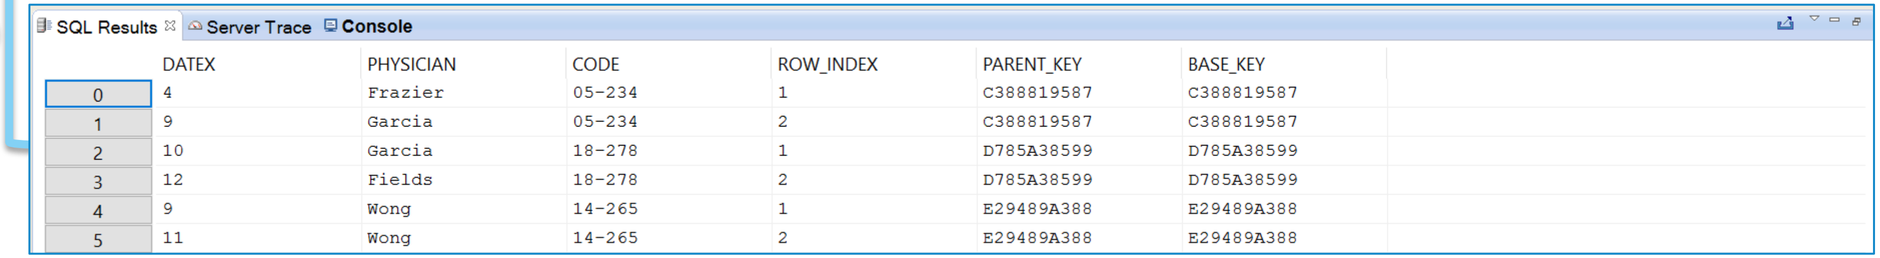

	DATEX	PHYSICIAN	CODE	ROW_INDEX	PARENT_KEY	BASE_KEY
0	4	Frazier	05-234	1	C388819587	C388819587
1	9	Garcia	05-234	2	C388819587	C388819587
2	10	Garcia	18-278	1	D785A38599	D785A38599
3	12	Fields	18-278	2	D785A38599	D785A38599
4	9	Wong	14-265	1	E29489A388	E29489A388
5	11	Wong	14-265	2	E29489A388	E29489A388

- Result Sets
  1. All data flattened in one table
  2. Parent table for VSAM cluster
  3. Child table for reoccurs

Parent/child to connect tables

Application Access

Access virtual tables from COBOL

```
      ******************************************************************01300001
       1000-OPEN.                                                       01310001
      ******************************************************************01320001
           DISPLAY 'OPENING CONNECTION TO DVS SERVER: ' DVCB-SSID.      01330001
           MOVE MESSAGE-LENGTH TO DVCB-MESSAGE-LENGTH.                  01340001
           MOVE FCT-OPEN TO DVCB-REQUEST-CODE.                          01350001
           CALL DSCLIENT USING DVCB SEND-AREA RECEIVE-AREA              01360001
                 ERROR-AREA SQLDA-AREA.                                 01370001
           PERFORM 5000-CHECK-RESPONSE.                                 01380001
      ******************************************************************01390001
       2000-SEND-SQL.                                                   01400001
      ******************************************************************01410001
           DISPLAY 'SENDING SQL TO DVS SERVER: ' DVCB-SSID.             01420001
           MOVE FCT-SEND TO DVCB-REQUEST-CODE.                          01430001
           MOVE SEND-LENGTH TO DVCB-SEND-LENGTH.                        01440001
           CALL DSCLIENT USING DVCB SEND-AREA RECEIVE-AREA              01450001
                 ERROR-AREA SQLDA-AREA.                                 01460001
           PERFORM 5000-CHECK-RESPONSE.                                 01470001
```

- Traditional COBOL coding
  1. Access data in place, VSAM, Sequential, IMS,etc.
  2. DSCLIENT embedded in COBOL code to access virtual tables

DVM code generator

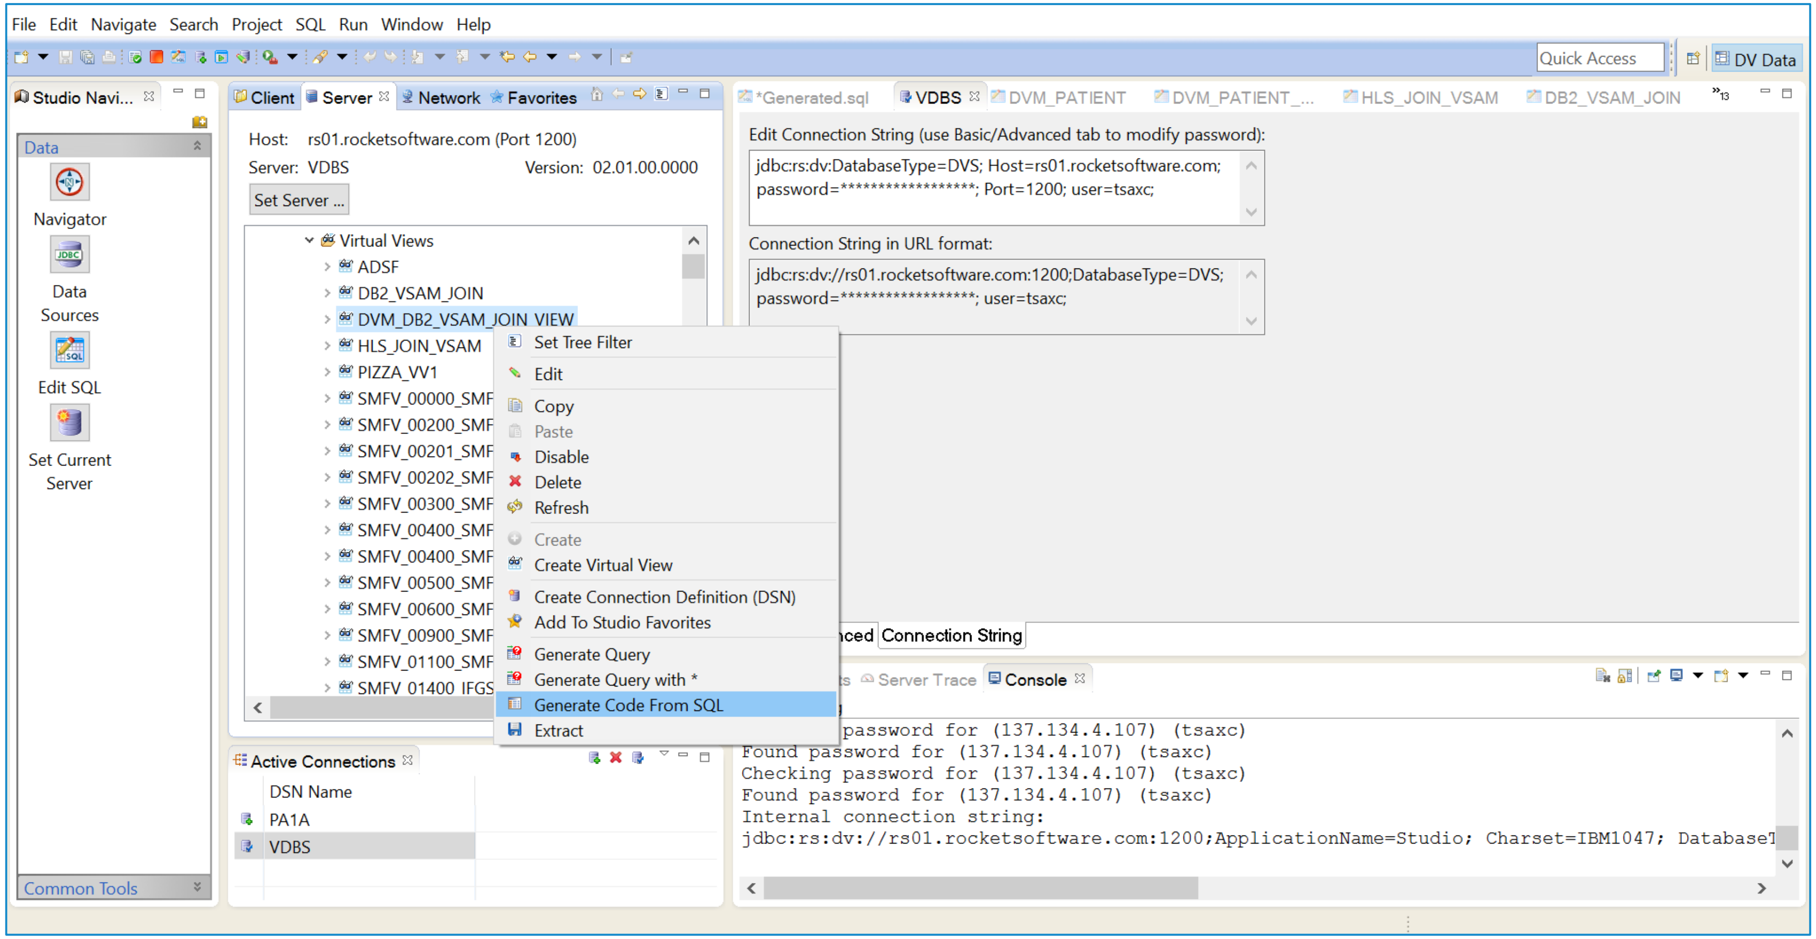

Host: rs01.rocketsoftware.com (Port 1200)

Server: VDBS Version: 02.01.00.0000

Edit Connection String (use Basic/Advanced tab to modify password):

jdbc:rs:dv:DatabaseType=DVS; Host=rs01.rocketsoftware.com; password=*****************; Port=1200; user=tsaxc;

Connection String in URL format:

jdbc:rs:dv://rs01.rocketsoftware.com:1200;DatabaseType=DVS; password=*****************; user=tsaxc;

```
password for (137.134.4.107) (tsaxc)
Found password for (137.134.4.107) (tsaxc)
Checking password for (137.134.4.107) (tsaxc)
Found password for (137.134.4.107) (tsaxc)
Internal connection string:
jdbc:rs:dv://rs01.rocketsoftware.com:1200;ApplicationName=Studio; Charset=IBM1047; DatabaseT
```

- Generating Code
  - Virtual table/view definition complete
  - Right-click on table/view and choose “Generate Code From SQL”

DVM code generator

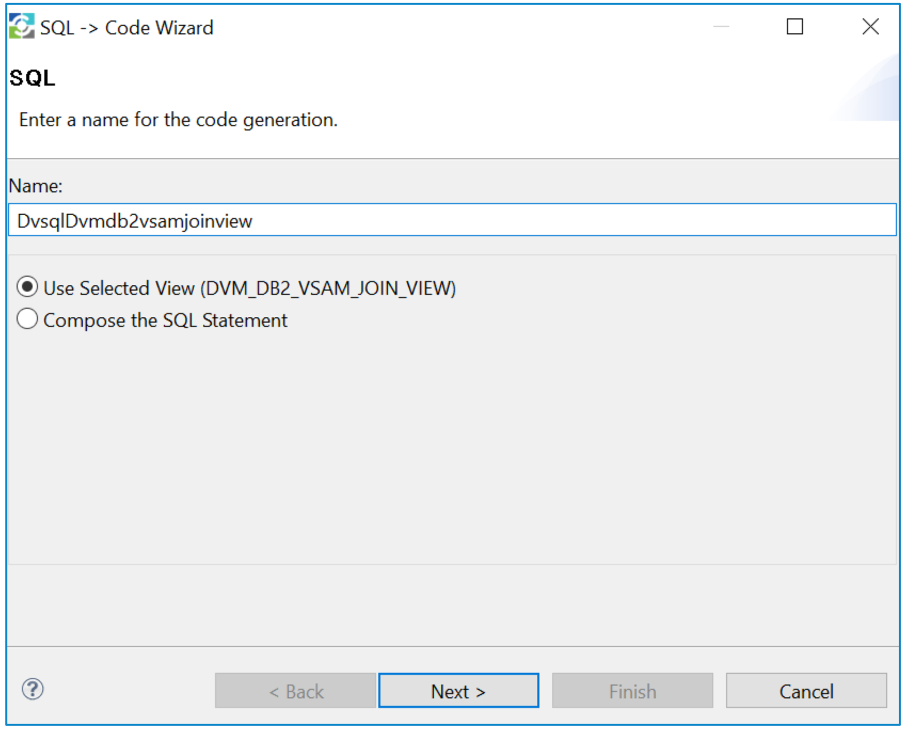
SQL -> Code Wizard

SQL

Enter a name for the code generation.

Name:

DvsqlDvmdb2vsamjoinview

● Use Selected View (DVM_DB2_VSAM_JOIN_VIEW)

○ Compose the SQL Statement

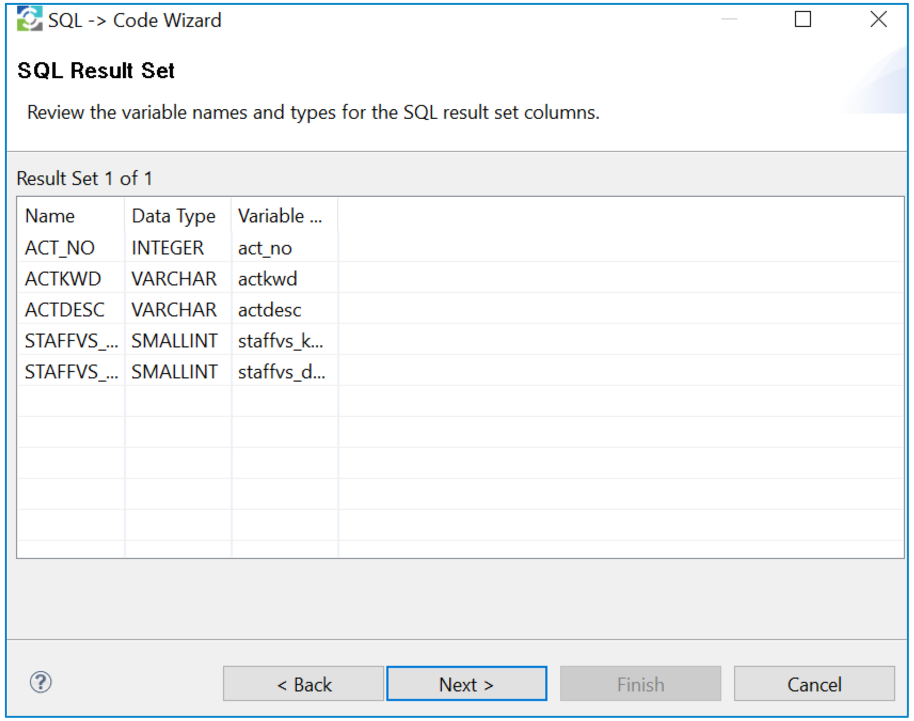
SQL -> Code Wizard

SQL Result Set

Review the variable names and types for the SQL result set columns.

Result Set 1 of 1

Name	Data Type	Variable ...
ACT_NO	INTEGER	act_no
ACTKWD	VARCHAR	actkwd
ACTDESC	VARCHAR	actdesc
STAFFVS_...	SMALLINT	staffvs_k...
STAFFVS_...	SMALLINT	staffvs_d...

- Generating Code
  - Input and target of code
  - Variables

Choose application type

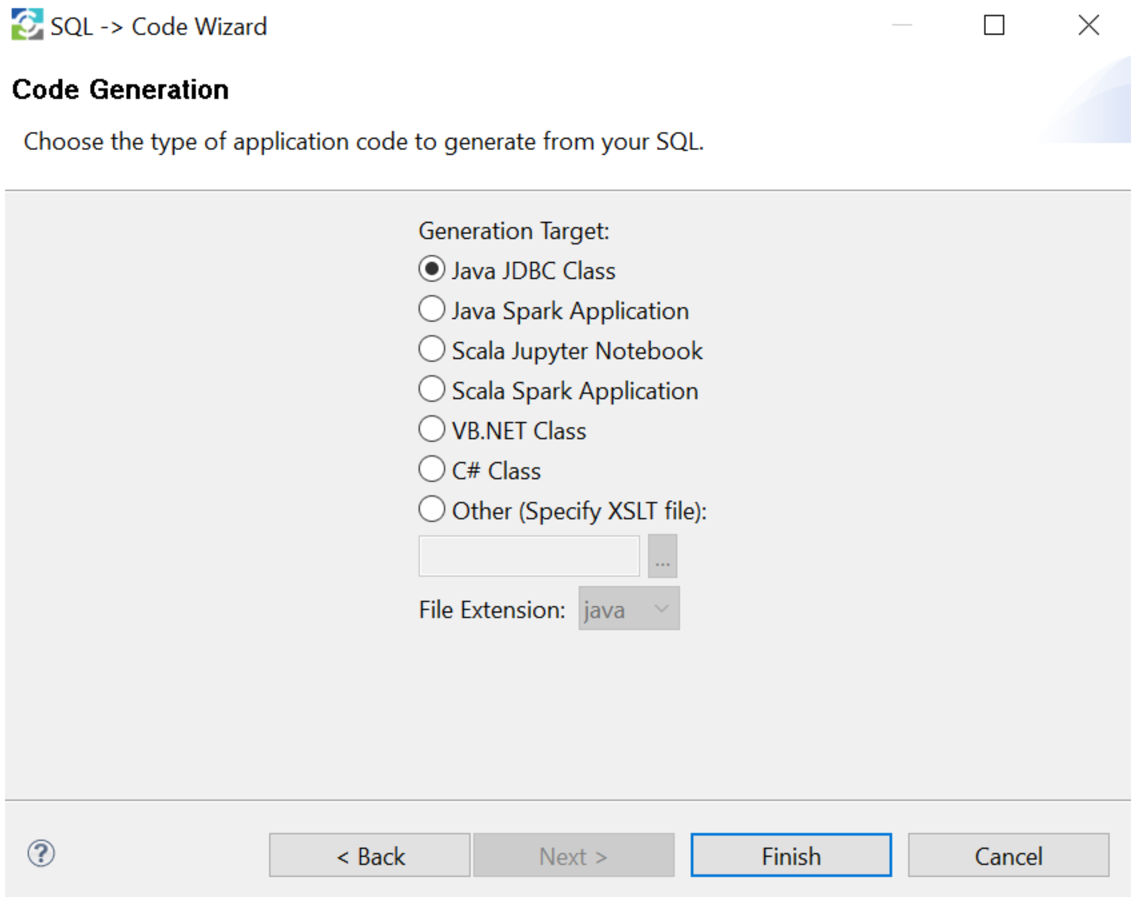

SQL -> Code Wizard

Code Generation

Choose the type of application code to generate from your SQL.

Generation Target:

- ● Java JDBC Class
- ○ Java Spark Application
- ○ Scala Jupyter Notebook
- ○ Scala Spark Application
- ○ VB.NET Class
- ○ C# Class
- ○ Other (Specify XSLT file):

File Extension: java

- Application code options

Generated Java code

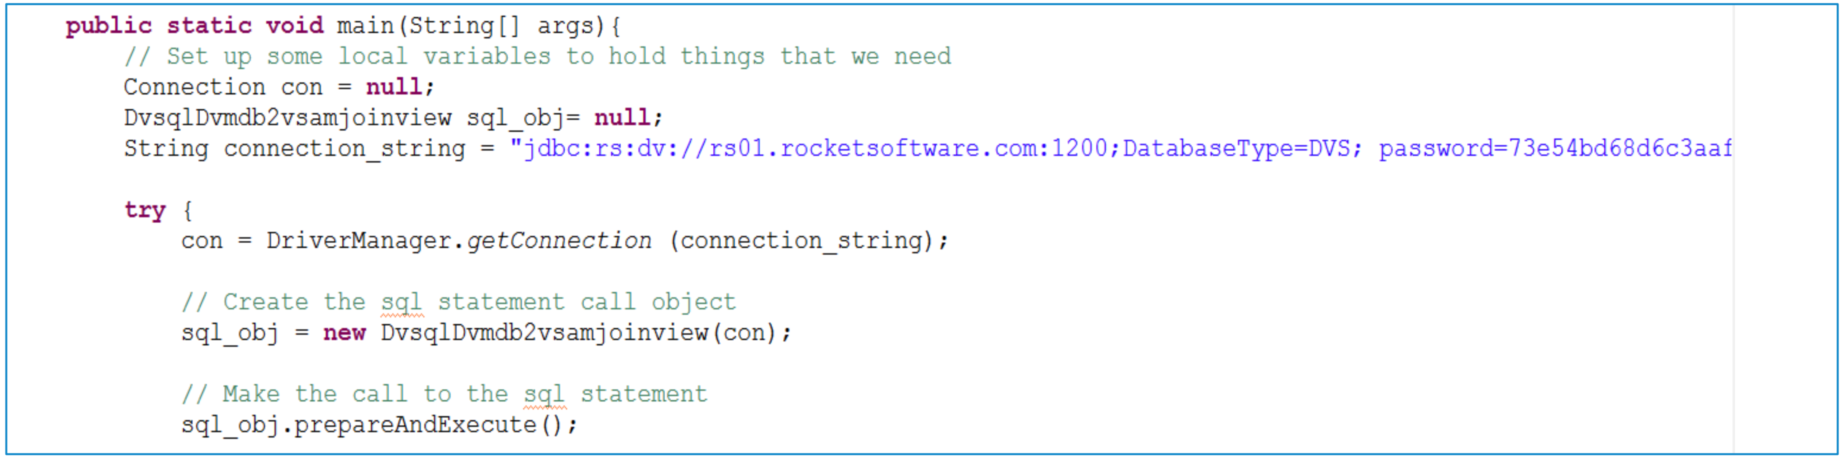

```
    public static void main(String[] args){
        // Set up some local variables to hold things that we need
        Connection con = null;
        DvsqlDvmdb2vsamjoinview sql_obj= null;
        String connection_string = "jdbc:rs:dv://rs01.rocketsoftware.com:1200;DatabaseType=DVS; password=73e54bd68d6c3aaf

        try {
            con = DriverManager.getConnection (connection_string);

            // Create the sql statement call object
            sql_obj = new DvsqlDvmdb2vsamjoinview(con);

            // Make the call to the sql statement
            sql_obj.prepareAndExecute();
```

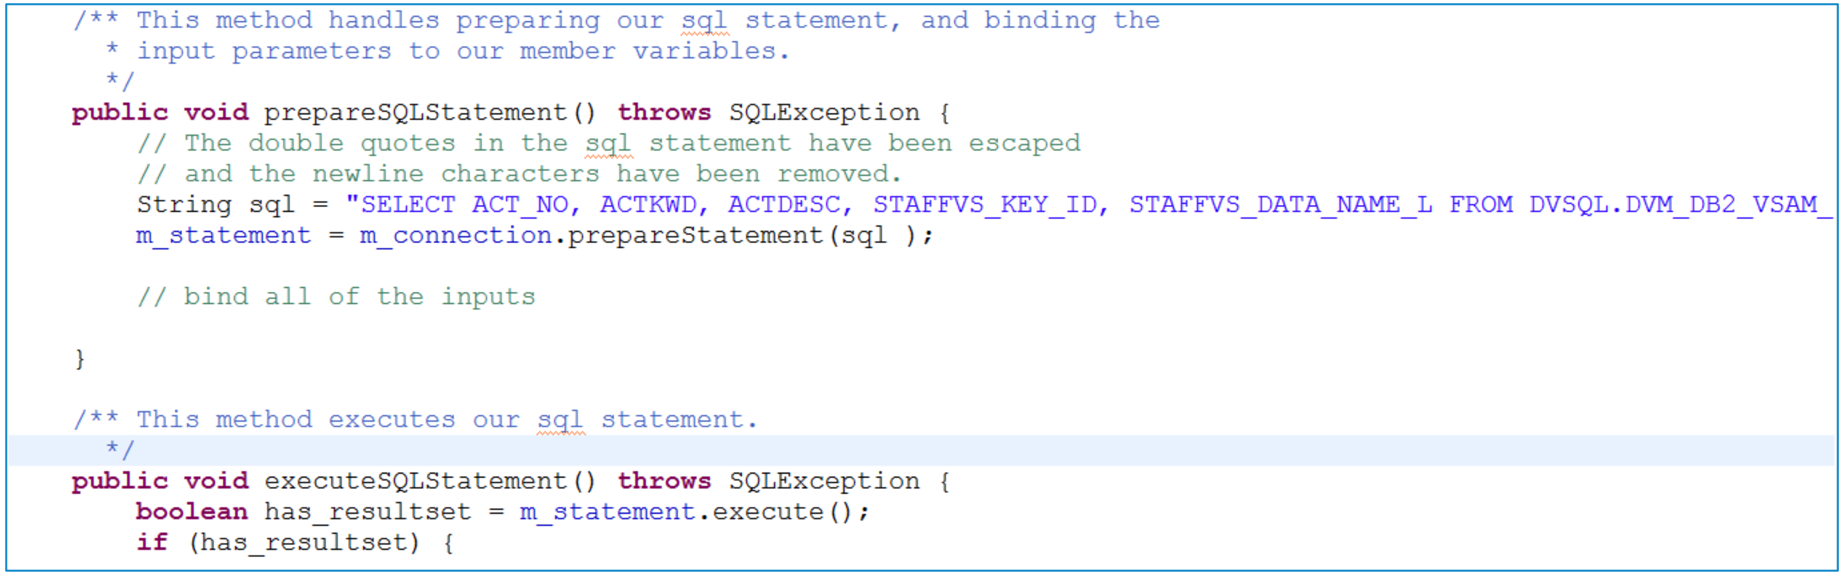

```
    /** This method handles preparing our sql statement, and binding the
      * input parameters to our member variables.
      */
    public void prepareSQLStatement() throws SQLException {
        // The double quotes in the sql statement have been escaped
        // and the newline characters have been removed.
        String sql = "SELECT ACT_NO, ACTKWD, ACTDESC, STAFFVS_KEY_ID, STAFFVS_DATA_NAME_L FROM DVSQL.DVM_DB2_VSAM_
        m_statement = m_connection.prepareStatement(sql );

        // bind all of the inputs

    }

    /** This method executes our sql statement.
      */
    public void executeSQLStatement() throws SQLException {
        boolean has_resultset = m_statement.execute();
        if (has_resultset) {
```

- Generated code is a complete working program
  - ➢ Samples show the connection to DVM and the SQL prepare and execute

ISPF interface

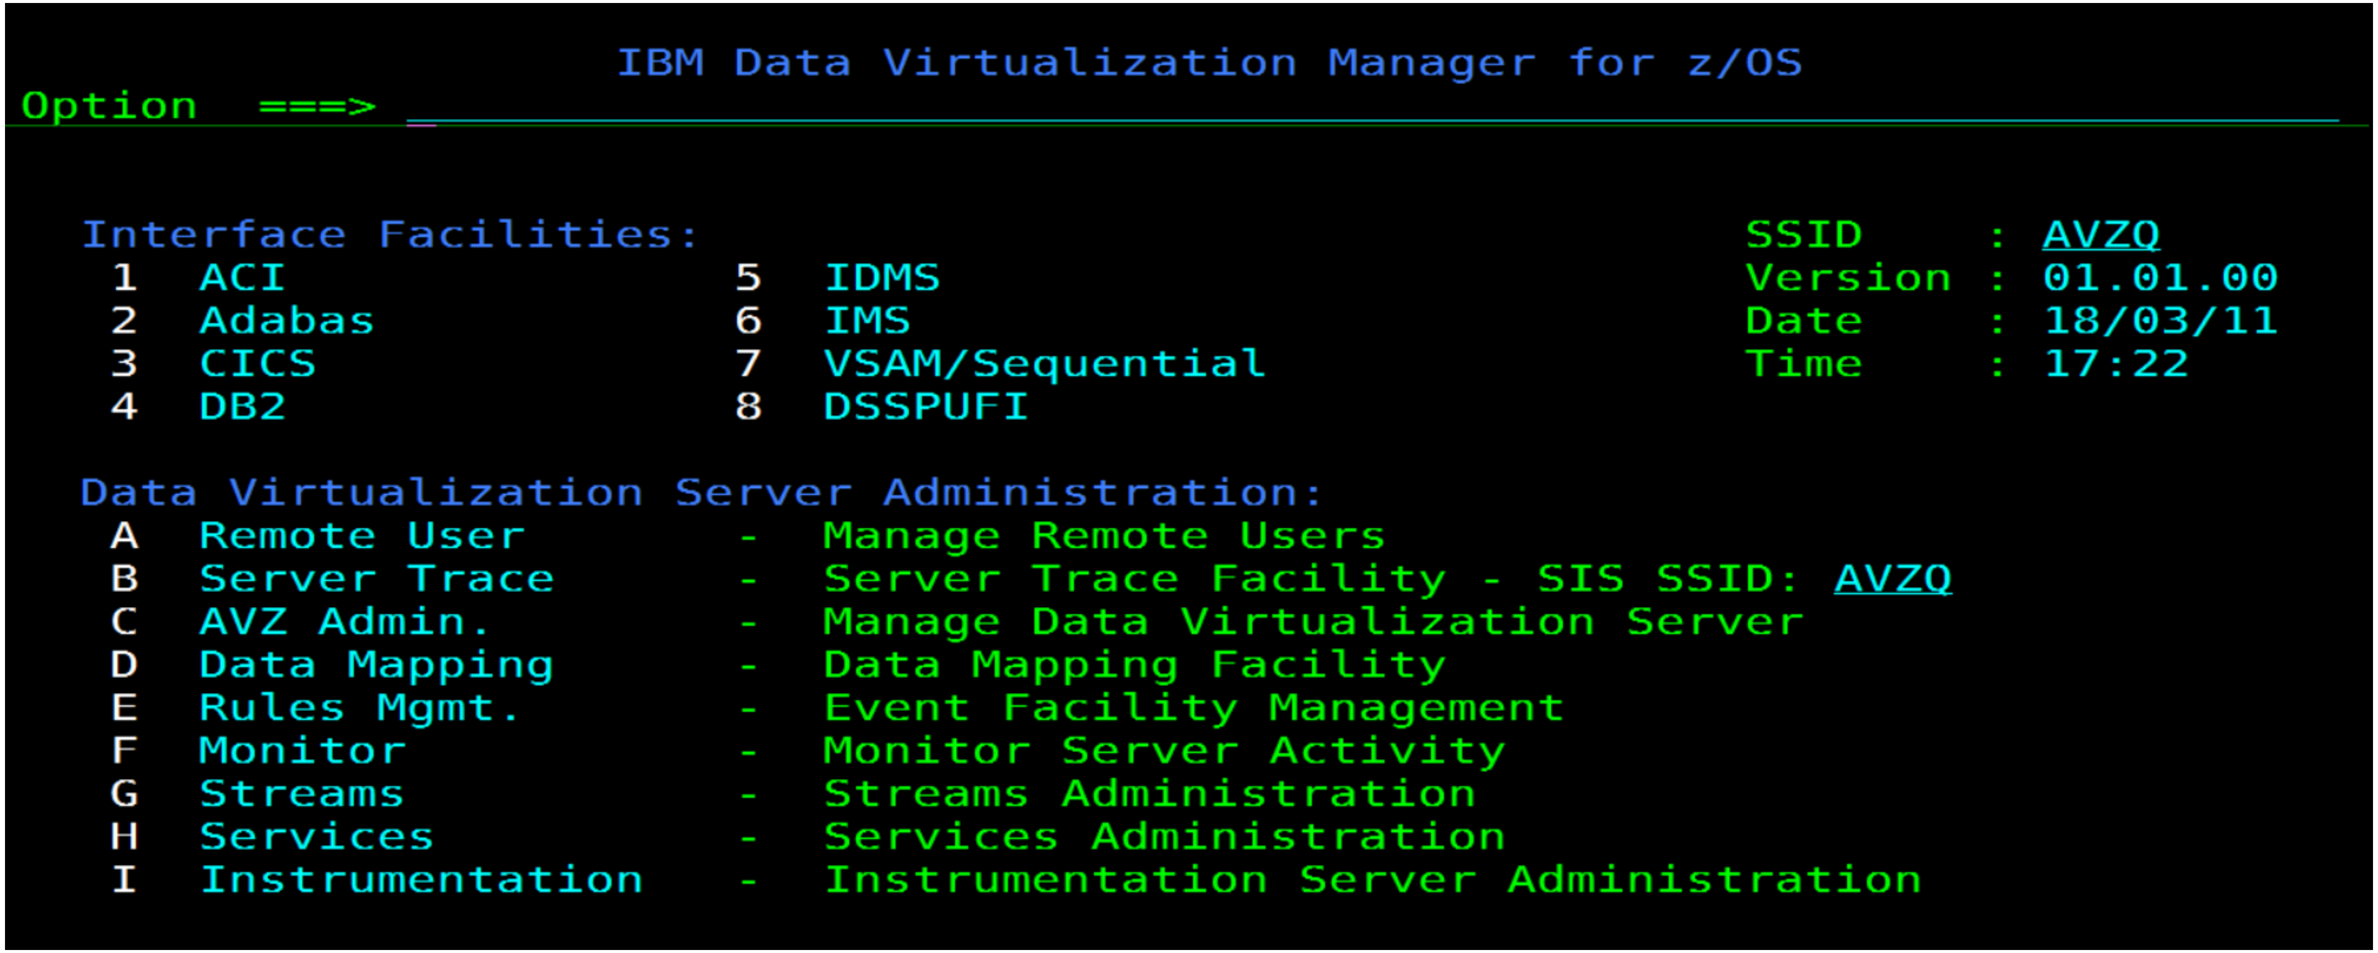

```
                    IBM Data Virtualization Manager for z/OS
Option   ===> 


  Interface Facilities:                                     SSID      : AVZQ
   1   ACI                   5   IDMS                       Version   : 01.01.00
   2   Adabas                6   IMS                        Date      : 18/03/11
   3   CICS                  7   VSAM/Sequential            Time      : 17:22
   4   DB2                   8   DSSPUFI

  Data Virtualization Server Administration:
   A   Remote User           -   Manage Remote Users
   B   Server Trace          -   Server Trace Facility - SIS SSID: AVZQ
   C   AVZ Admin.            -   Manage Data Virtualization Server
   D   Data Mapping          -   Data Mapping Facility
   E   Rules Mgmt.           -   Event Facility Management
   F   Monitor               -   Monitor Server Activity
   G   Streams               -   Streams Administration
   H   Services              -   Services Administration
   I   Instrumentation       -   Instrumentation Server Administration
```

Trace browse

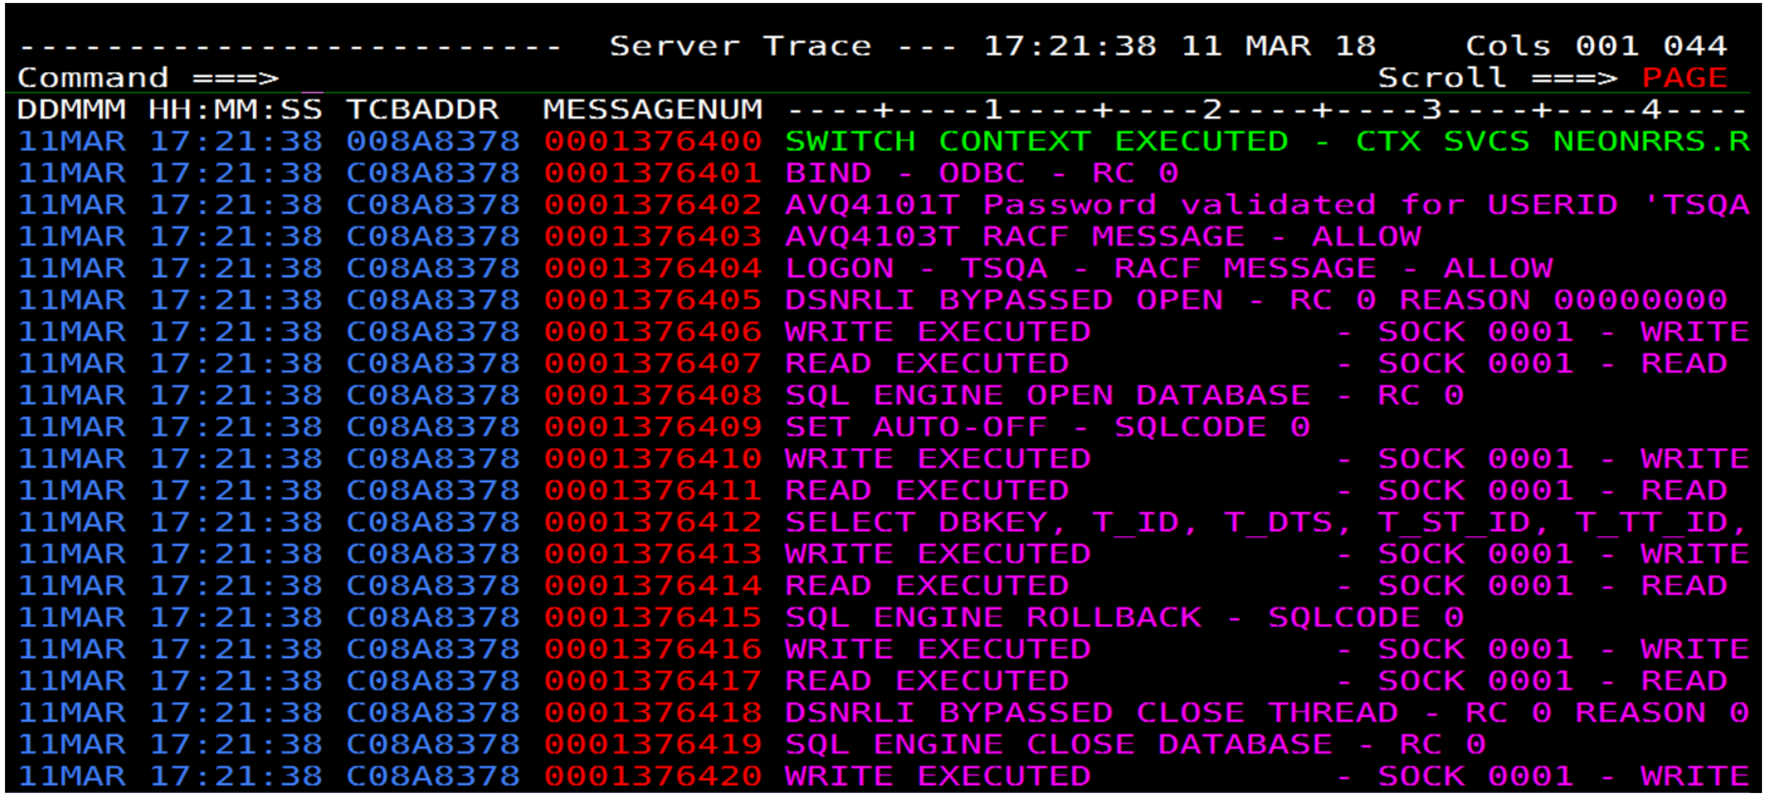

```
-------------------------   Server Trace --- 17:21:38 11 MAR 18    Cols 001 044
Command ===> _                                             Scroll ===> PAGE
DDMMM HH:MM:SS TCBADDR  MESSAGENUM ----+----1----+----2----+----3----+----4----
11MAR 17:21:38 008A8378 0001376400 SWITCH CONTEXT EXECUTED - CTX SVCS NEONRRS.R
11MAR 17:21:38 C08A8378 0001376401 BIND - ODBC - RC 0
11MAR 17:21:38 C08A8378 0001376402 AVQ4101T Password validated for USERID 'TSQA
11MAR 17:21:38 C08A8378 0001376403 AVQ4103T RACF MESSAGE - ALLOW
11MAR 17:21:38 C08A8378 0001376404 LOGON - TSQA - RACF MESSAGE - ALLOW
11MAR 17:21:38 C08A8378 0001376405 DSNRLI BYPASSED OPEN - RC 0 REASON 00000000
11MAR 17:21:38 C08A8378 0001376406 WRITE EXECUTED           - SOCK 0001 - WRITE
11MAR 17:21:38 C08A8378 0001376407 READ EXECUTED            - SOCK 0001 - READ
11MAR 17:21:38 C08A8378 0001376408 SQL ENGINE OPEN DATABASE - RC 0
11MAR 17:21:38 C08A8378 0001376409 SET AUTO-OFF - SQLCODE 0
11MAR 17:21:38 C08A8378 0001376410 WRITE EXECUTED           - SOCK 0001 - WRITE
11MAR 17:21:38 C08A8378 0001376411 READ EXECUTED            - SOCK 0001 - READ
11MAR 17:21:38 C08A8378 0001376412 SELECT DBKEY, T_ID, T_DTS, T_ST_ID, T_TT_ID,
11MAR 17:21:38 C08A8378 0001376413 WRITE EXECUTED           - SOCK 0001 - WRITE
11MAR 17:21:38 C08A8378 0001376414 READ EXECUTED            - SOCK 0001 - READ
11MAR 17:21:38 C08A8378 0001376415 SQL ENGINE ROLLBACK - SQLCODE 0
11MAR 17:21:38 C08A8378 0001376416 WRITE EXECUTED           - SOCK 0001 - WRITE
11MAR 17:21:38 C08A8378 0001376417 READ EXECUTED            - SOCK 0001 - READ
11MAR 17:21:38 C08A8378 0001376418 DSNRLI BYPASSED CLOSE THREAD - RC 0 REASON 0
11MAR 17:21:38 C08A8378 0001376419 SQL ENGINE CLOSE DATABASE - RC 0
11MAR 17:21:38 C08A8378 0001376420 WRITE EXECUTED           - SOCK 0001 - WRITE
```

Pink – executed on zIIP
Green – executed on GPP

TCBADDR:
C* - zIIP and GPP same speed
D* - zIIP faster than GPP

Line commands:
P/PP – print
S/SS – print with zoom

Primary commands:
DISPLAY
PROFILE

A Modern Option to Data Federation

- InfoSphere Classic Federation Server for z/OS - provides integration of non-relational z/OS data with distributed tools and applications through SQL-driven access to z/OS data sources

IBM® Data Virtualization Manager for z/OS® - allows organizations to virtualize z/OS data with other enterprise data sources in real-time to provide comprehensive information.

API support limited to SQL	Support for any API – SQL, NoSQL, REST, SOAP…
No support for zIIP specialty engine	Up to 99% zIIP eligible for low TCO
No machine specific software versions	Support for 4 optimized machine specific versions (z14, z13, EC12, and z196)
Performance - does not support parallelism	High Performance - built in MapReduce parallelism
No support for DRDA	Full DRDA Support
No industry data model support	Pre-built metadata maps for ISO 8583 and 20022

ANY Data for ANY Application

Simple

Get transactional access, no data movement

Open to all Apps

Modern APIs enable access

Secure

Avoid risk by reducing moving data off Z Systems

Fast

Exploits Z architecture, including parallelism and in-memory processing

Cost Effective

Keeps Z costs down with up to 99% zIIP offload

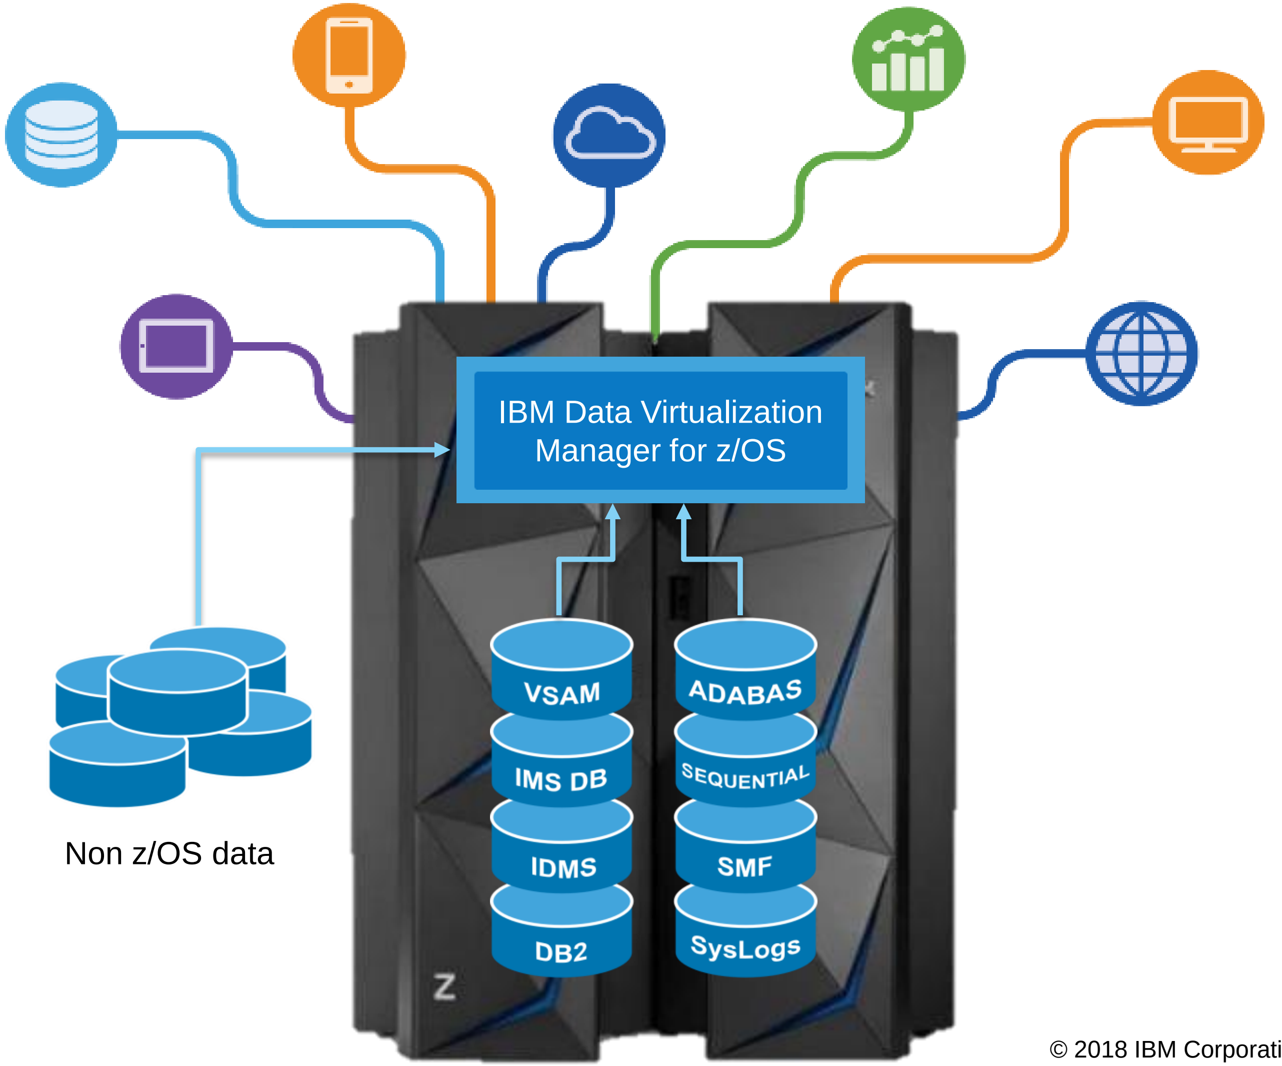

Where to Go for More Information

IBM Announcement
http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/4/897/ENUS217-404/index.html&lang=en&request_locale=en

Product documentation
http://www-304.ibm.com/support/docview.wss?uid=swg27020910#ibmdvmz-lib

YouTube "IBM Data Virtualization Manager for z/OS" channel
https://www.youtube.com/channel/UCtbd_4oHoH-uKDYgSSRL7SA

Be sure to subscribe for new videos and to hit “Like” button