Compiler Forecast: Cloudy with a Chance of Savings

Cloud Computing for your z/OS compilers

www.cloudcompiling.com

CLOUD COMPILING

Speaker Bio

- Charles Mills is the Chief Development Officer of Cloud Compiling, LLC
- Mills was the founder of Firesign Computer Company, the developers of Outbound, which became the number three mainframe file transfer product in the market. Mills sold his company to Allen Systems Group (ASG)
- Cloud Compiling develops a line of cloud or virtual compilers for z/OS

Agenda

- What do we mean by the phrase “cloud” or “virtual” compiler?
- What about … ?
- Development Issues
- Why would a customer use a cloud compiler?

A slide for the lawyers…

Any Compiler is a Black Box

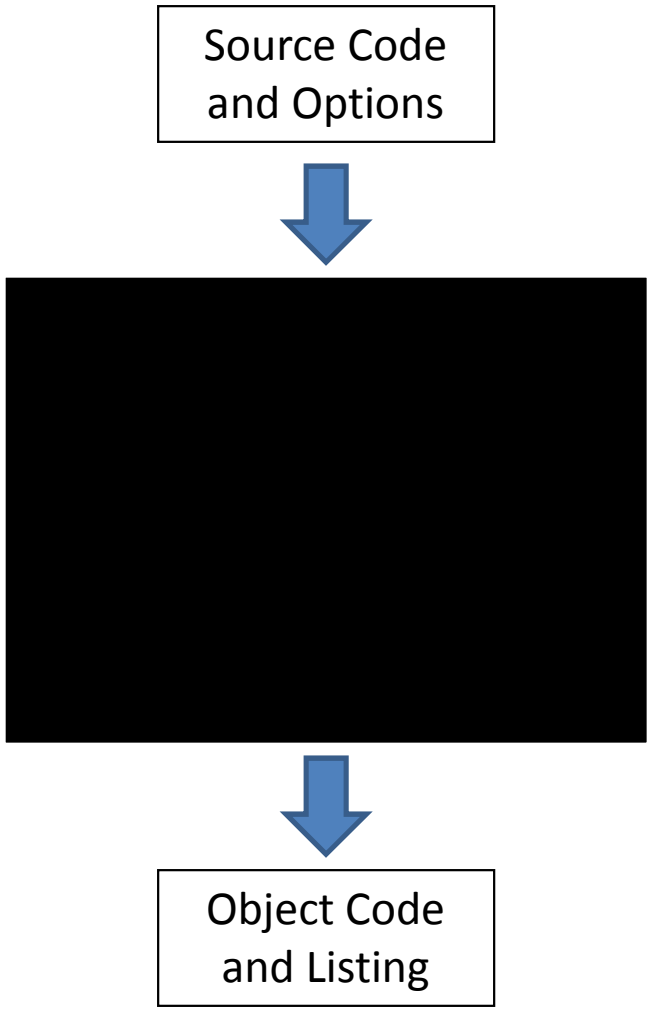

How a Cloud Compiler Works

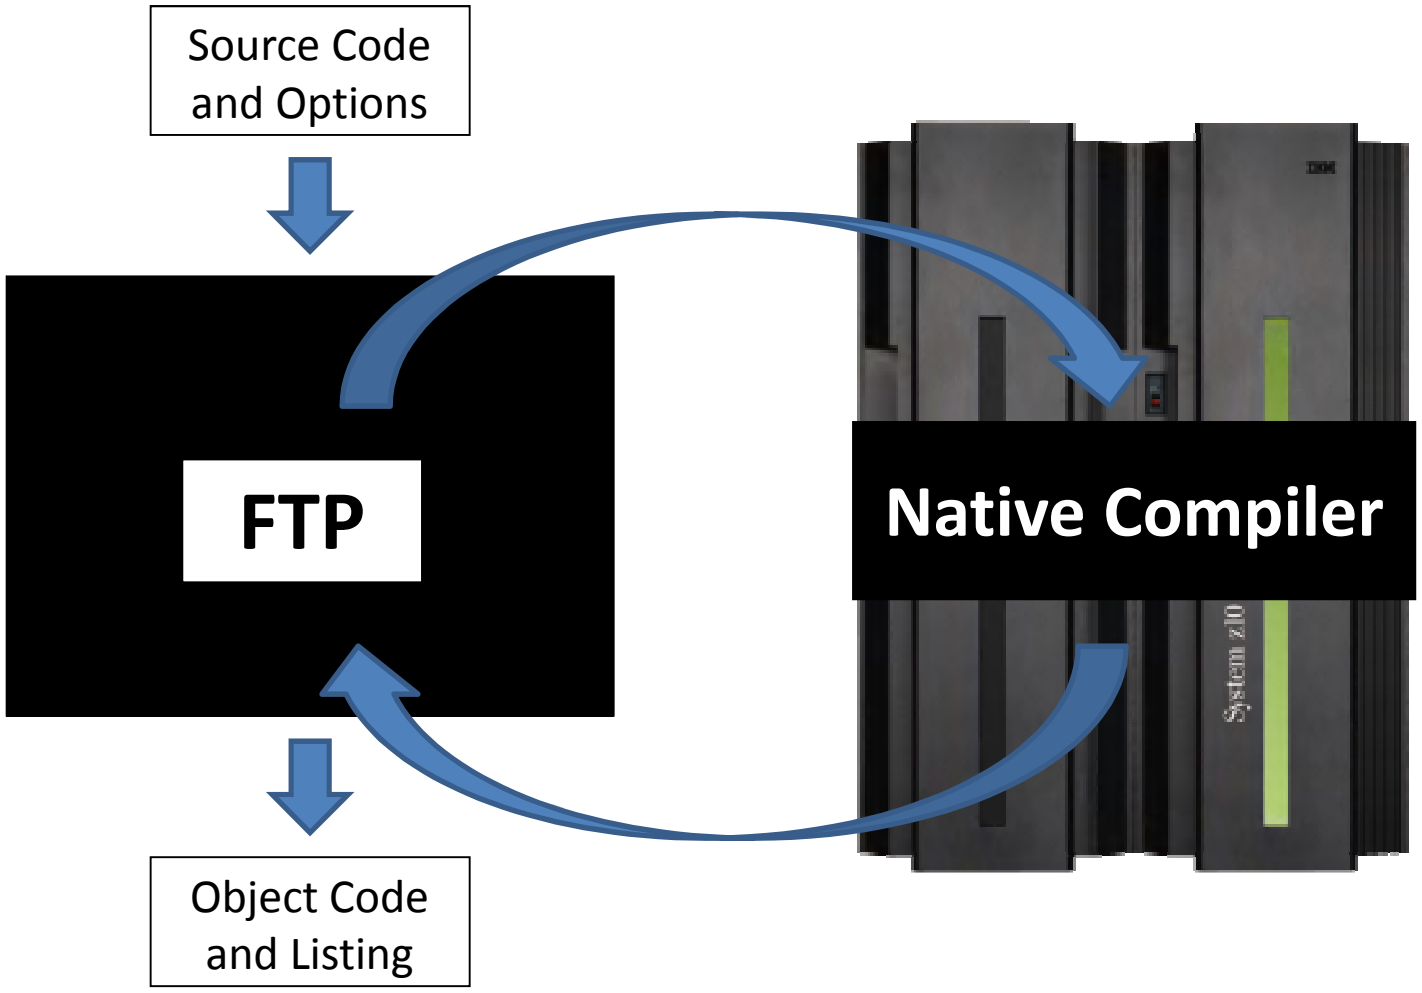

Transparency

```
//COMPILE   EXEC PGM=TCCENTCB,
//    PARM='OBJ,ADATA,SZ(4096)'
//STEPLIB   DD   DSN=TCC.LOADLIB,DISP=
//SYSIN     DD   DSN=MYSOURCE...
//SYSLIB    DD   DSN=MYLIB...
//          DD   DSN=OTHERLIB...
//SYSLIN    DD   DISP=(NEW,PASS),...
//SYSPRINT  DD   SYSOUT=*
//TCCPARMS  DD   DSN=TCC.PARM.FILE...
//TCCPRINT  DD   SYSOUT=*
```

How Our Cloud Compiler Works

- Analyzes environment: DD's, PARM=, etc.
- Reads through source code
- FTPs source code to target system
- Builds a new compile job and uses FTP to submit
- After job completes FTPs object code and listing back
- Note does *not* "move the job from one JES to another" or anything like that
  - Does not require JESPLEX or close coupling
  - Only requires an FTP (TCP/IP) link
  - Supports mixed JES2/JES3 environment

A Totally New Job

```
//XBC001D  JOB  (),'CHARLES MILLS',MSGCLASS=H,REGION=0M
//* Generated by Cloud Compiler PLIOPTV2 V3R0.0 at 14:21:53 on 7/15/2010
//*  on behalf of job/step RDC001A(JOB05072)/CLOUDCMP on node TESTJES
//PLIOPTV2 EXEC PGM=IEL0AA,COND=(0,NE),REGION=4M,
//  PARM=('SIZE(1000K),MARGINS(2,72),X,NOBJ,DECK')
//STEPLIB  DD   DISP=SHR,DSN=PLI230.PLICOMP
//SYSIN    DD   DISP=(OLD,DELETE),
//  DSN=XBC001.TCCFILES.XBC001D.FHM92JZA.SYSIN
//SYSLIB   DD   DISP=(OLD,DELETE),
//  DSN=XBC001.TCCFILES.XBC001D.FHM92JZA.SYSLIB
//SYSPRINT DD   BLKSIZE=0,DISP=(NEW,CATLG),
//  DSN=XBC001.TCCFILES.XBC001D.FHM92JZA.SYSPRINT,DSORG=PS,LRECL=133,
//  RECFM=FBA,SPACE=(CYL,(1,1,0))
//SYSPUNCH DD   BLKSIZE=0,DISP=(NEW,CATLG),
//  DSN=XBC001.TCCFILES.XBC001D.FHM92JZA.SYSPUNCH,DSORG=PS,LRECL=80,
//  RECFM=FB,SPACE=(TRK,(5,10,0))
//SYSUT1   DD   SPACE=(CYL,(2,2,0)),UNIT=SYSDA
```

Two Kinds of Cloud

Open Cloud

Your Organization | Cloud Host

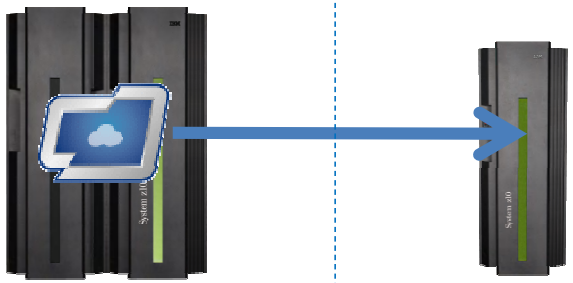

Private Cloud

Your Organization

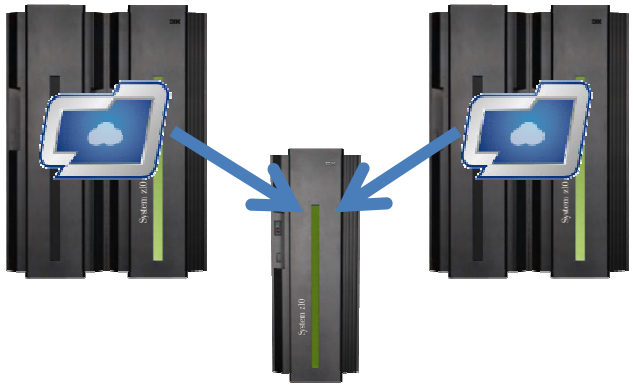

What about …

- What about COPY and INCLUDE?
- What about the _____ compiler option?
- What about the line going down?
- What about IBM? Is this legal?
- What about the link editor?
- What about CA-Endevor and Xpediter?
- What about the DB2 and CICS pre-compilers?
- What about the load on our network?
- What about compiler customization?
- What about security?

What about COPY and INCLUDE?

- Cloud compiler parses source code COPY or INCLUDE (as appropriate to language)
- FTPs relevant members and builds remote SYSLIB
- No source code stored on compile machine
  - *No synchronization issues*
- Handles default and “DD name” format
  - COPY member OF ddname
- Handles nested COPYs

What about the ____ compiler option?

- Short answer: “no problem”
- Longer answer
  - Enterprise COBOL supports 59 options
    - ADATA, ADV, ARITH, AWO, BUFSIZE, …
  - We care about 13 of them
    - ADATA means need to process SYSADATA
    - DECK means need to process SYSPUNCH
    - LIB means need to scan source code for COPY
    - Etc.
  - ADV, ARITH, AWO, BUFSIZE, etc. mean nothing to us
  - We pass them all to the compiler unmodified

What about the line going down?

- Extensive diagnostics, FTP “deadman,” etc.
- Customizable retry count
- Ability to define multiple compile servers
  - Automated fall-back
- For single datacenter clouds, if network down programmers probably dead in the water anyway
- Enable Safe-Cloud™ feature if desired
  - Falls back to IBM compiler installed on same machine
  - Legal to leave installed and not pay so long as don’t use
  - Safe-Cloud puts out audit message and you owe IBM for the month

What about IBM? Is this legal?

- Private cloud
  - Perfectly legal to route all of your compiles to one machine
- Open cloud
  - Our licenses with IBM permit compiles as a service
  - No different than if your programmers used us as a service bureau
- Major customers

What about the link editor?

- Link editor/binder licensed with z/OS, not compilers
- Link edit/bind in normal way after compile
- Link editor/binder does not know/care where object code came from
  - Remember the black box analogy

What about CA-Endevor and Xpediter?

- They don’t care
- Cloud compiler “looks just like” native compiler
  - Remember the black box analogy
  - They call an entry point and expect certain datasets to appear
- Some issues solved along the way
  - Example: hard-coded compiler entry point name

What about DB2 and CICS pre-compilers and DB2 bind?

- Licensed with DB2 and CICS, not compiler
- Run before or after cloud compiler just like native compiler
  - Remember the black box analogy
- Co-compiler requires DB2 or CICS installed on compile machine
  - Same version a good idea!
- You can always use the pre-compiler

What about the load on our network?

- Compile data volumes are surprisingly low
  - Two to three megabytes is a large compile
  - Like one medium-sized digital photograph
- Benchmark COBOL compile: 5489 lines
  - SYSPRINT: 21,431 lines or 2.8 MB
    - MAP option accounts for over half of that
  - 3 milliseconds over FICON Express8
  - 11 milliseconds over 2 Mb FICON
  - 1.3 seconds over ESCON
  - Fifteen seconds over T-1

What about compiler customization?

- "We have unique corporate compiler default options – will we lose them?"
  - Customize exactly as now but on compile server mainframe
- "We have two programming groups with different default options"
  - Several ways to handle – documented in our manuals
- Open cloud – multiple customers on our machine
  - We have devised a way to handle multiple customizations

What about security?

- No source code stored on compile machine
- Exists on compile machine for only a few seconds
  - `//SYSIN DD DISP=(OLD,DELETE),...`
  - Protected by RACF and ENQ'ed by z/OS
  - Cloud compiler uses FTP to delete if compile job totally fails
  - Unpredictable name like XCC001.TCCFILES.XCC001K.EORVA12U.SYSIN
- PassTickets
- Suggest defining userid with no TSO and limited dataset access
- Private Cloud
  - Just as secure as any other kind of compile
- Open Cloud
  - Secure technologies such as VPN, SSL, PassTickets, etc.

Interesting technical problems we solved

- No JCL changes
- Concatenated PDSes
- Subsystem PDSes
- FTP performance on PDSes
- PassTickets
- PARM=EXIT(…)

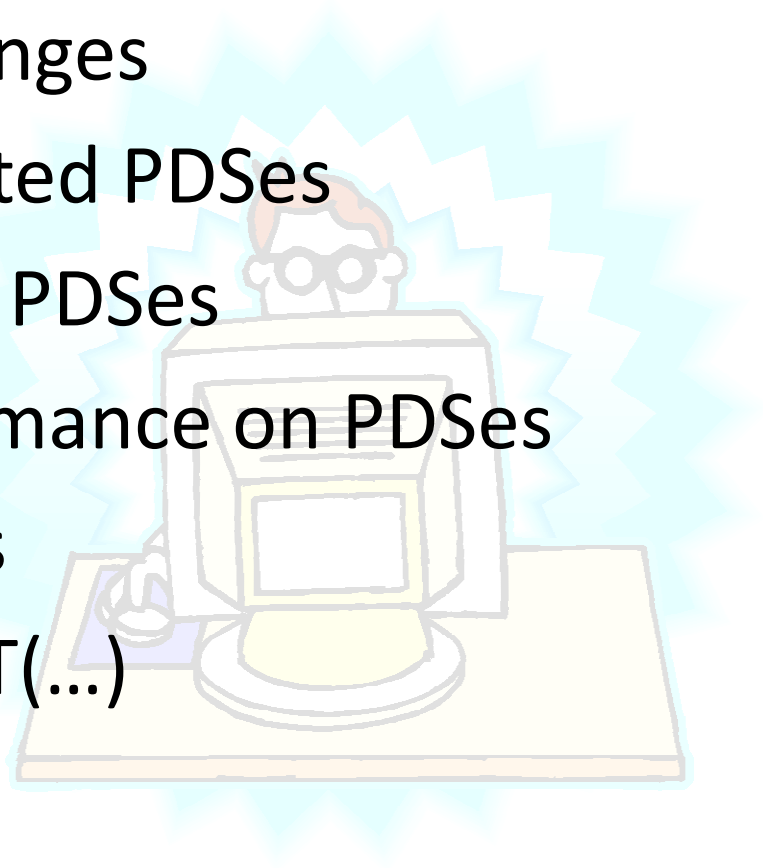

No JCL Changes

- Initial design goal: *minimal* JCL changes
  - "Change a couple of PROCs and you're all set"
  - Entry point and load library names
  - Add DDs for our unique parameter and log files
  - FTP does not set LRECL and RECFM so add to DDs
- Second customer an outsourcer
  - Not a clue where all the PROCs were
- Requirement became no JCL changes
  - Support for native compiler name as alias
  - Parameter and log files dynamically allocated
  - Pre-create certain datasets to set LRECL and RECFM

PDSes

- How you transmit PDS members with FTP

```
CD    THE.REMOTE.PDS
LCD   MY.LOCAL.PDS
PUT   MYMEMBR1
PUT   MYMEMBR2   NEWNAME
LCD   MY.OTHER.PDS
PUT   MYMEMBR3
```

Concatenated PDSes

- JCL:

```
//SYSLIB DD DSN=our.first.copy.library,DISP=SHR
       DD DSN=some.other.library,DISP=SHR
       DD DSN=another.copy.library,DISP=SHR
```

- FTP:

```
LCD ? ; Can't LCD DDname or concatenated PDSes
PUT member
```

- Solution:
  - C library fldata() returns true DSname (unless zFS!)
    - Additional code for zFS simulated PDSes
  - Or with assembler can use BLDL and RDJFCB with ARL

Subsystem PDSes

- JCL:

```
//SYSLIB DD DSN=some.libr.file,SUBSYS=LAM
```

- FTP:

```
LCD some.libr.file    ; Not really a PDS
PUT member            ; Does not work!
```

- Solution:
  - Dynamically allocate a new DD name to dataset *and member* with SUBSYS=LAM or PAM
  - **`PUT //DD:ddname  memname`**

Poor FTP performance on PDS members

- FTP uses SVC 99 to allocate a new DD name for every PDS member
- DYNALLOC plus OPEN about ten times as expensive as BPAM FIND
  - Adds up if hundreds of COPYs (IMS!)
- Solution:
  - Dynamically invoke IEBCOPY to unload specified members to single sequential dataset
  - FTP to remote compile machine
  - In compile job use EXEC PGM=IEBCOPY to rebuild
  - Negative: IEBCOPY requires APF authorization

PARM=EXIT(…)

- Original design goal: support all commonly-used compiler options
- Who uses EXIT?
- Answer: IBM File Manager (File-AID replacement)
- Problem: File Manager code highly dependent on undocumented specifics of native compiler implementations of EXIT
- We finally got it working

PassTickets

- The problems
  - FTP requires a password
  - Don’t want to transmit passwords over the network
  - Don’t want to – or auditors won’t let us – store passwords
  - Concern that access to one application may give access to others
- The solution …

PassTickets

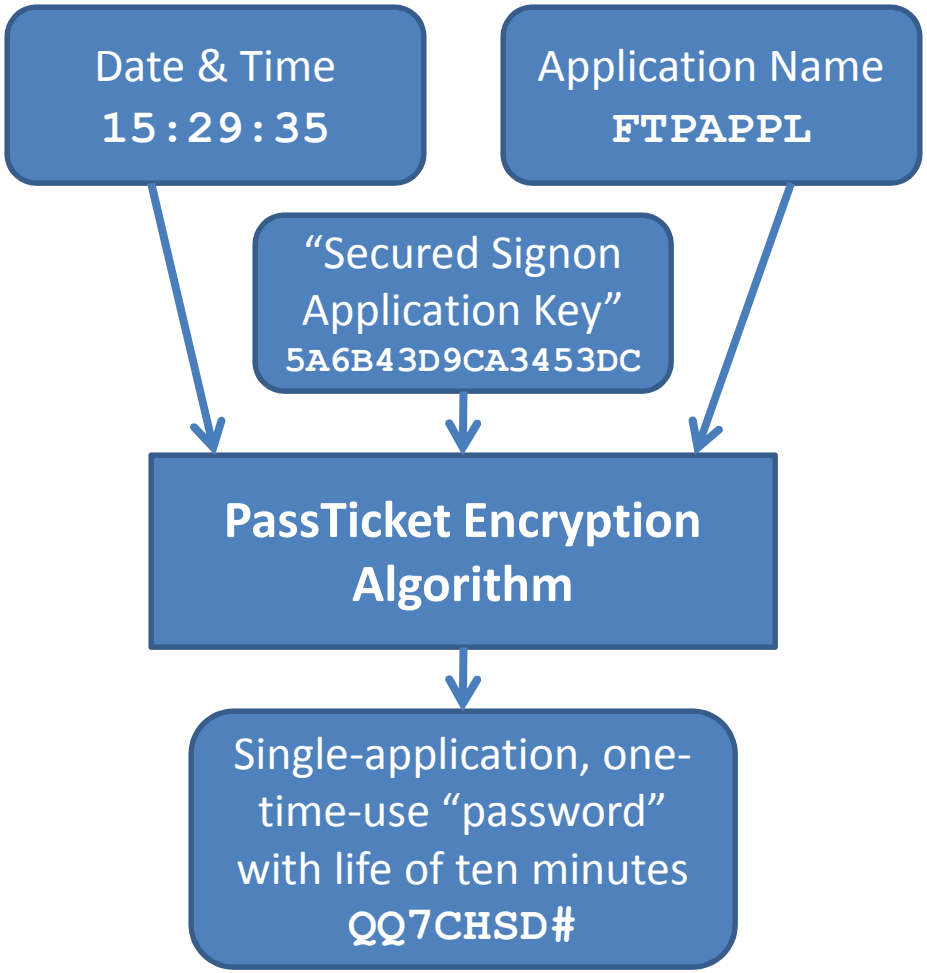

- Exactly like a password … but
  - No exposure in transmitting
  - No reason to store
  - Access to one application (FTP)
- Requires mainframe clocks set within a few minutes of each other (usually UTC)
- Program must be APF-authorized to generate
- Was an absolute bear for us to figure out
- But now have "cookbook" in our manuals
- Yes, supported also by ACF2 and TopSecret

This all sounds complicated

- Don’t meant to give the impression that using cloud compiling is complicated
- All of the things discussed happen automatically under the hood

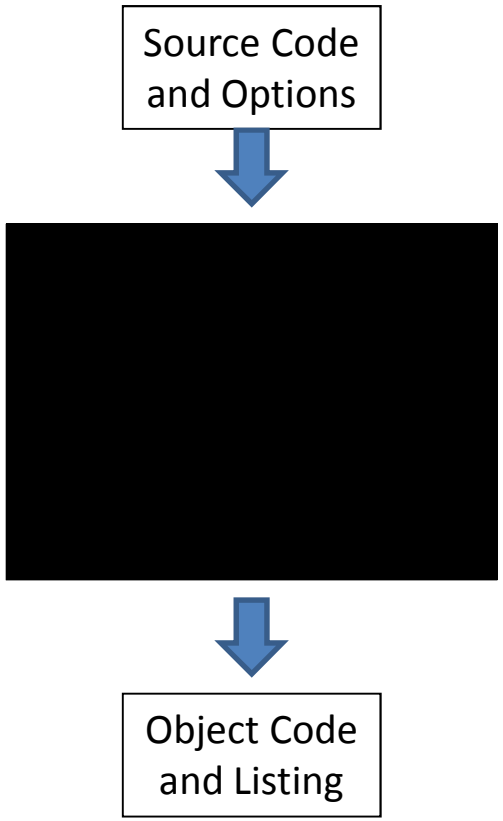

Why would a customer use a cloud compiler?

Total $20,428 per Month

Why would a customer use a cloud compiler?

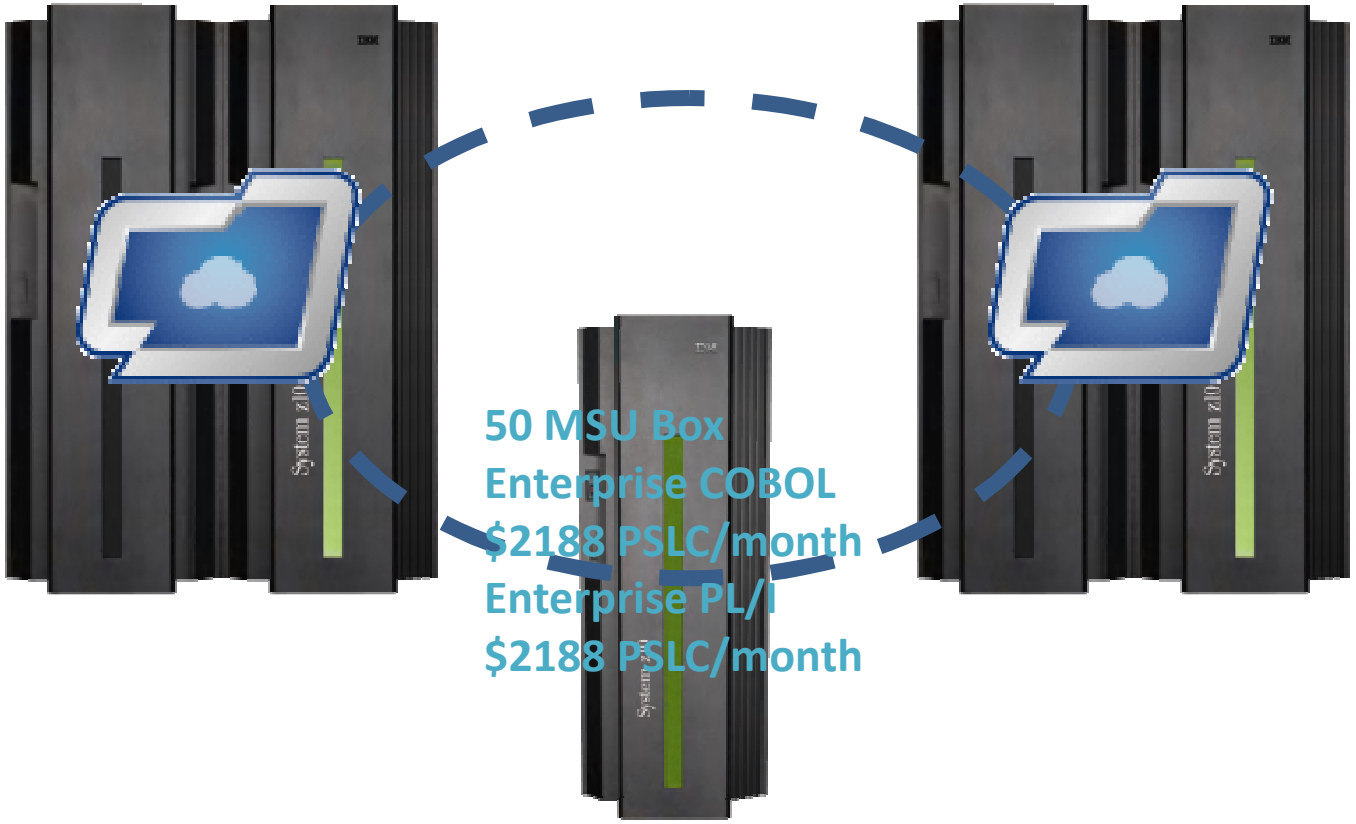

PSLC $4376 + Cloud Compiling $8026/month (50% of savings)
Total $12,402/month

A few words from a customer

Questions?

Or ask me off-line: Booth 423 or Charles.Mills@CloudCompiling.com

Summary

- What do we mean by “cloud” compiler?
- What about … ?
- Development Issues
- Why would a customer use a cloud compiler?

- Thank you for attending