

The Power to Stream z IT Operational Data to the Analytic Engine of Your Choice

Domenico D'Alterio
IBM

November 2018
Session OK



Agenda

- Business challenges
- IBM Common Data Provider for z Systems
 - Overview
 - Product configuration
 - User scenarios
 - Cross-product integrations
 - Roadmap & Strategy

Business challenges

The Digital Economy is forcing businesses to transform

Business transformation effects on IT



- **Explosion in transaction growth**

- *driven by mobility and the Internet of Things*



- **Analytics is moving to real time**

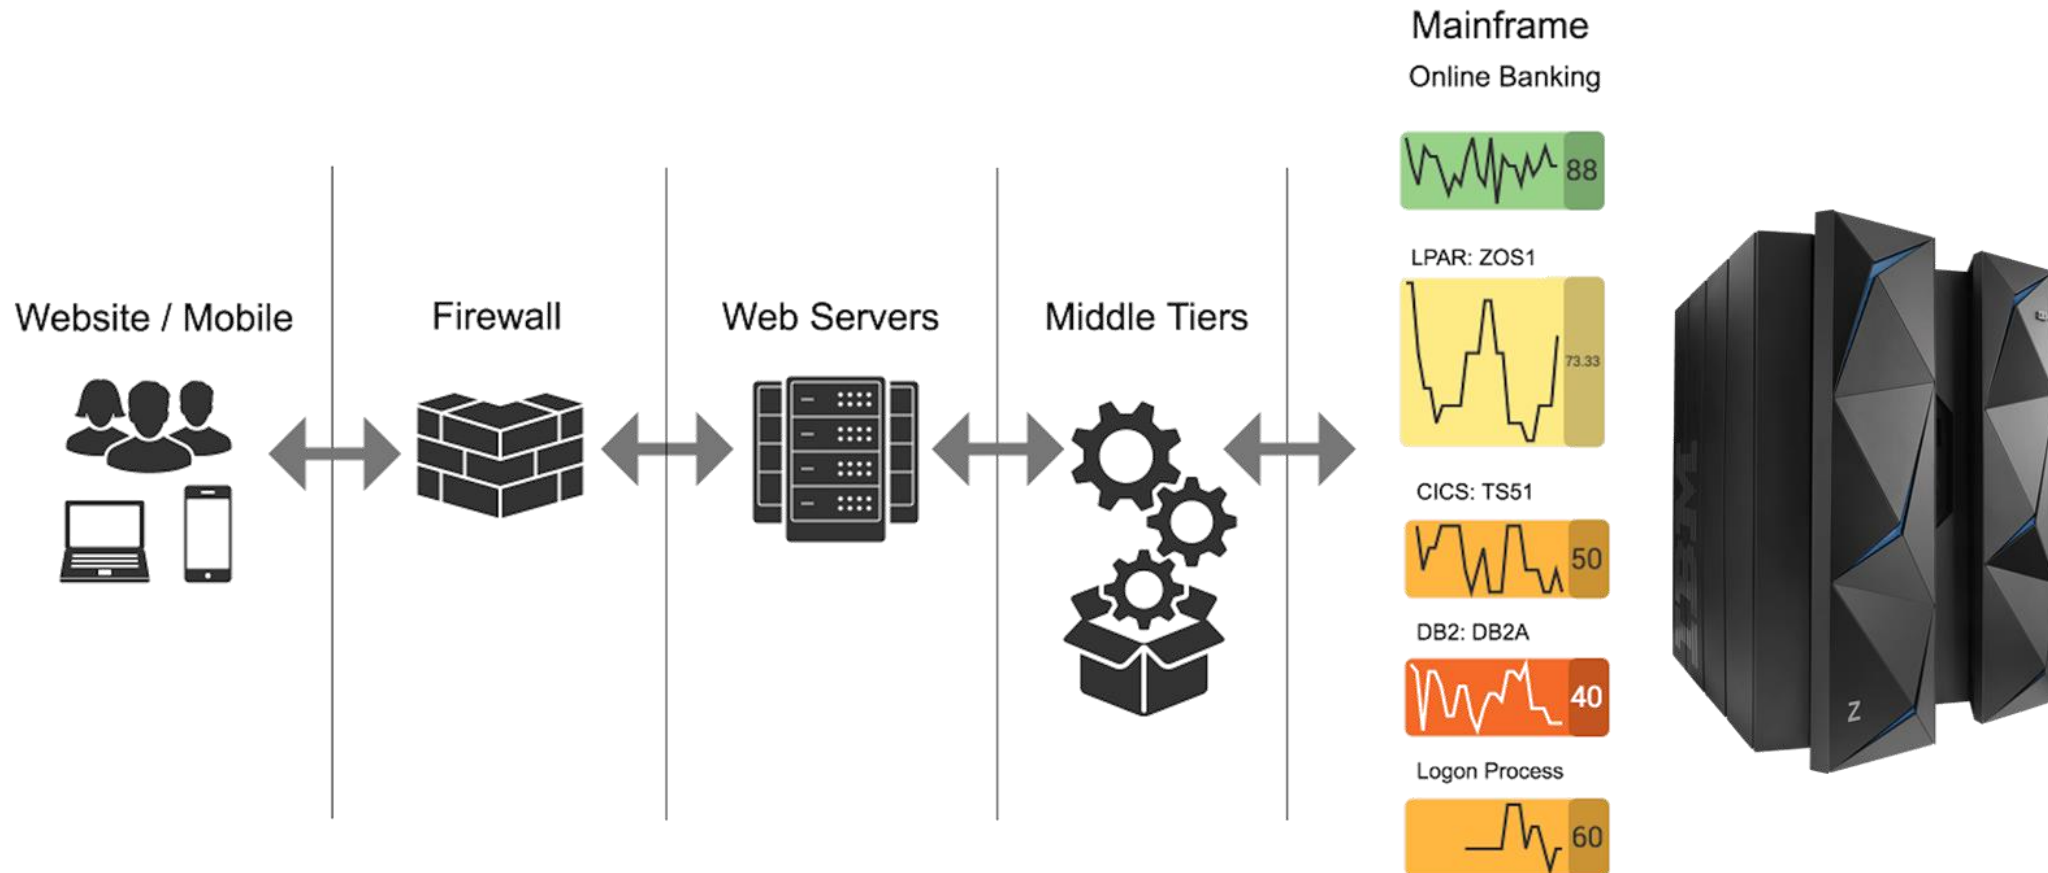
- *to capture new opportunities at the point of impact*




- **IT- driven business agility**

- *for delivering service, security, and efficiency*

View of Today's Hybrid IT Operations




The mainframe is literally just a big black box to our lines of business, security, and compliance teams. **Getting access to Z data** will eliminate this blind spot.



Z IT Operations challenge

Multiple Data Sources – increasing number of consumers



Same data requested by different consumers
Different data to be sent to the same consumers

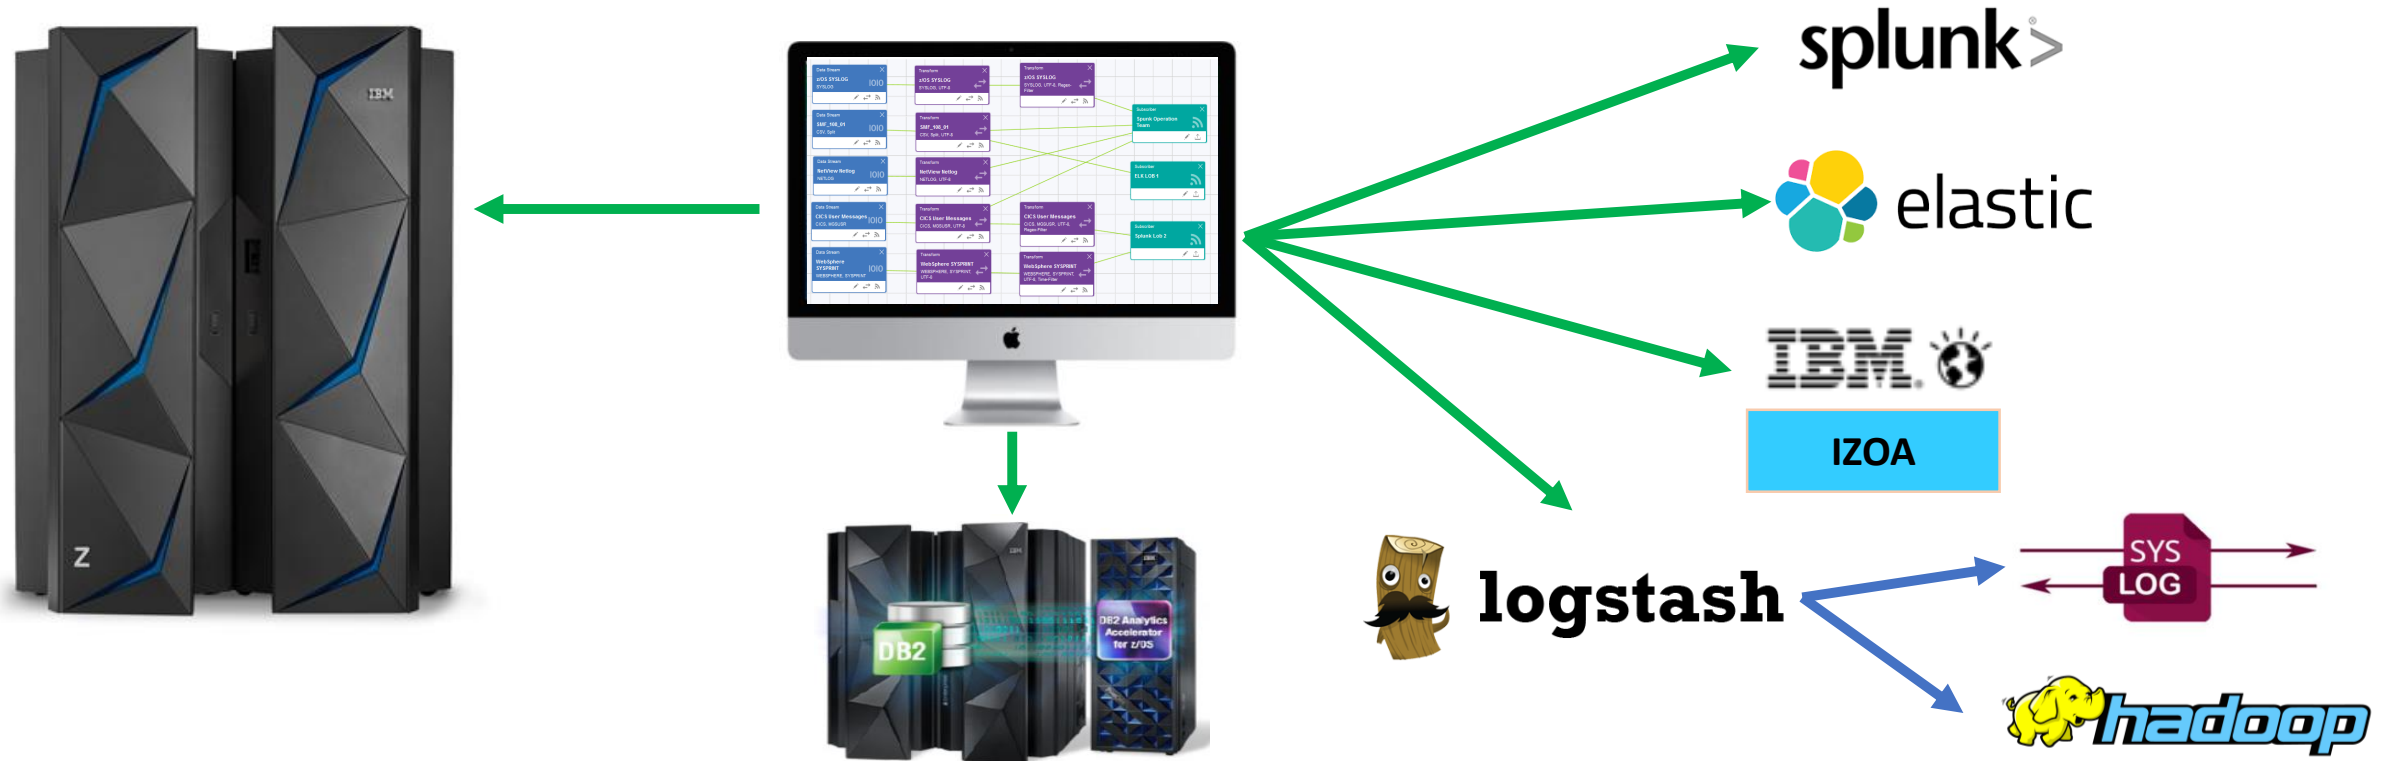
IBM Common Data Provider for z Systems

Overview

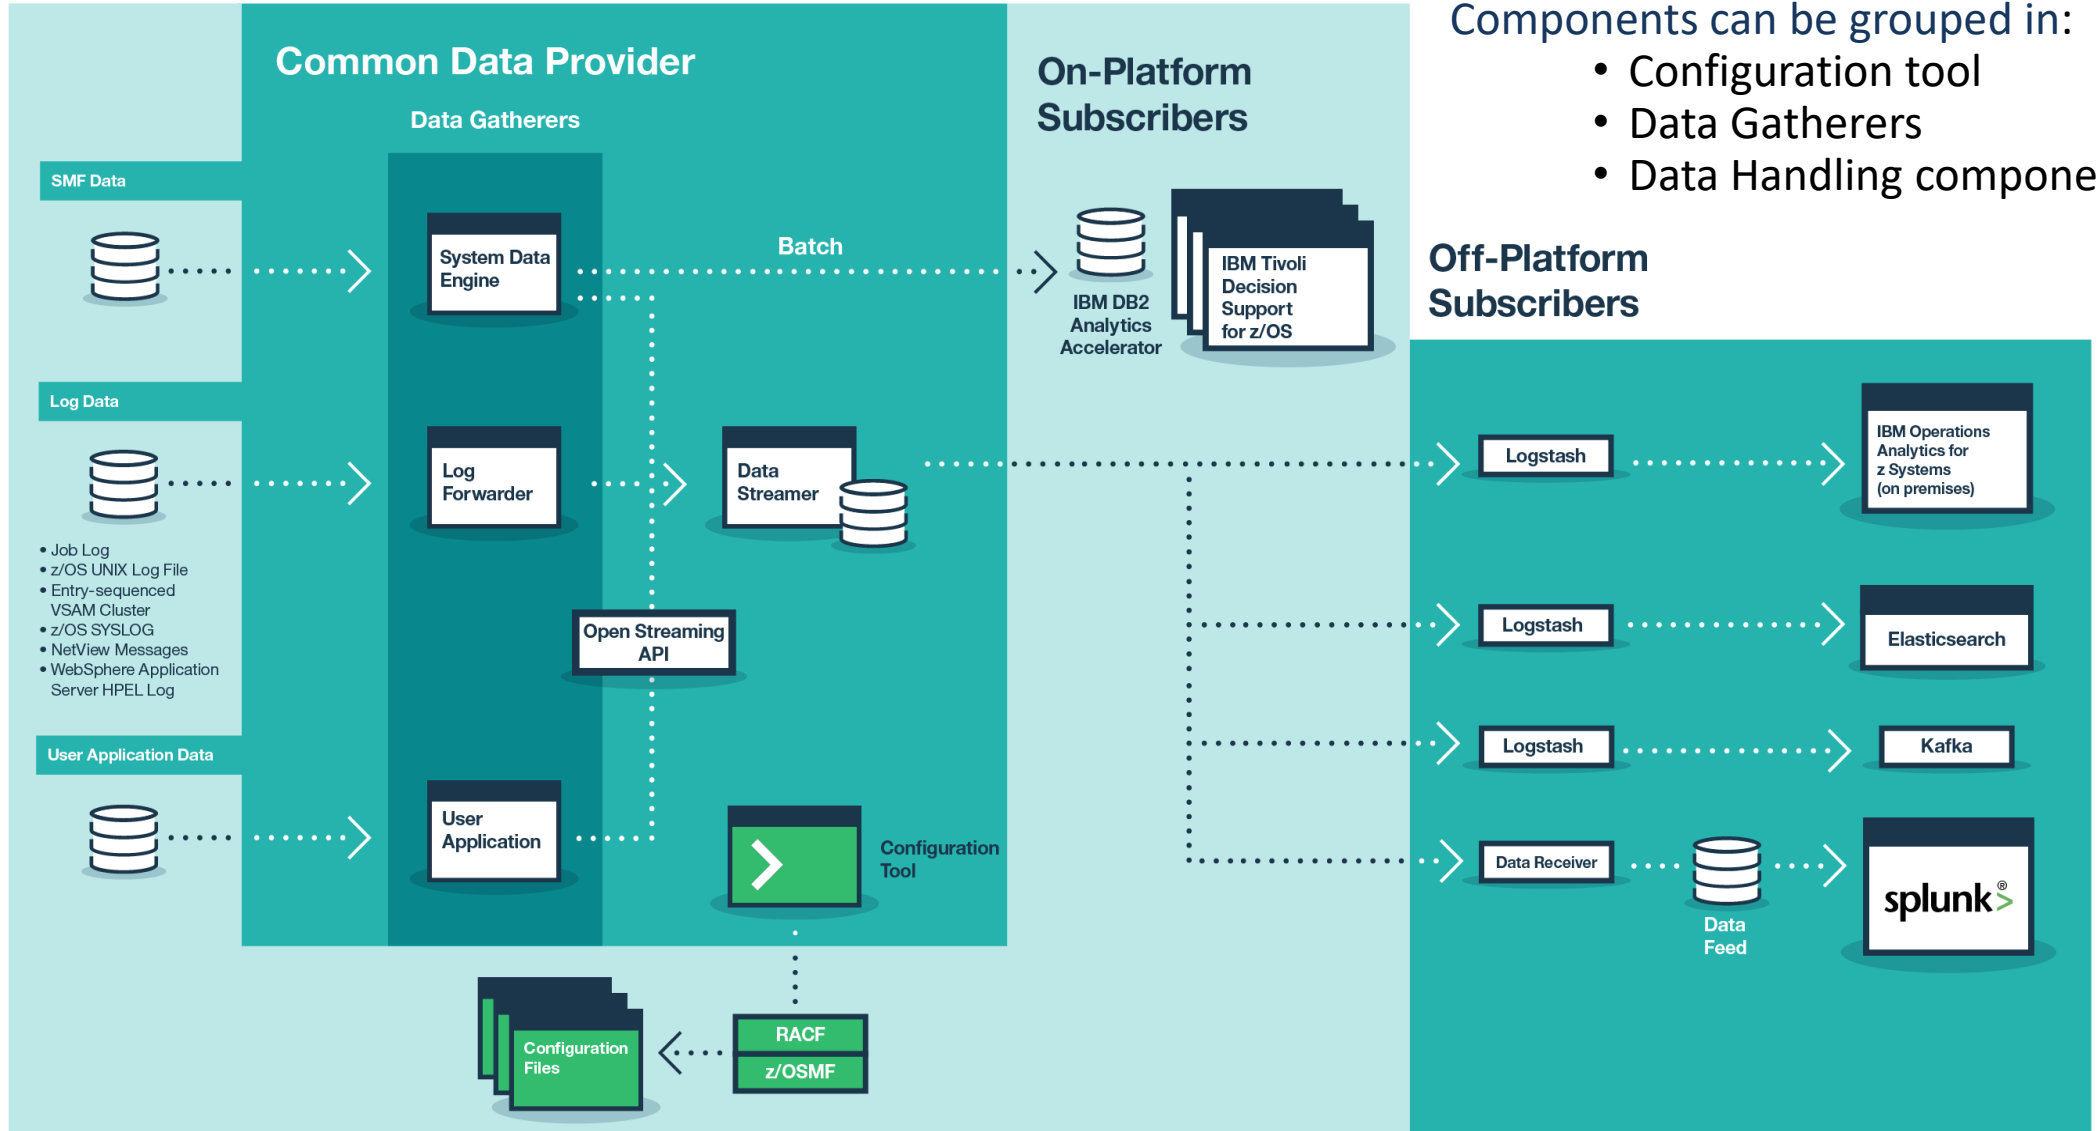
IBM Common Data Provider for z Systems

Near real time collection of structured and unstructured IT operational data available to your analytics solution

IBM Common Data Provider for z Systems (CDPz) enables users to gather IBM Z IT operational data through a **single interface**, providing structured and unstructured data in **near real time** to a variety of analytics solutions. Data can be provided both on and off platform in a consistent, **consumable** format.



IBM Common Data Provider for z Systems Architecture



- Components can be grouped in:
- Configuration tool
 - Data Gatherers
 - Data Handling components

Web Configuration Tool

Common Data Provider can be configured through a user-friendly Web Configuration Tool



- The **configuration** of the **data sources** to be leveraged, the **transformations** to be applied on the data collected and **target consumers** of the data can be performed all through a single user friendly **Web User Interface** (Plug-in for **z/OS Management Facility**):
- **Security:** Host-based policy controls subscribers and data sources. Policy can be secured by RACF for total control of data and subscribers



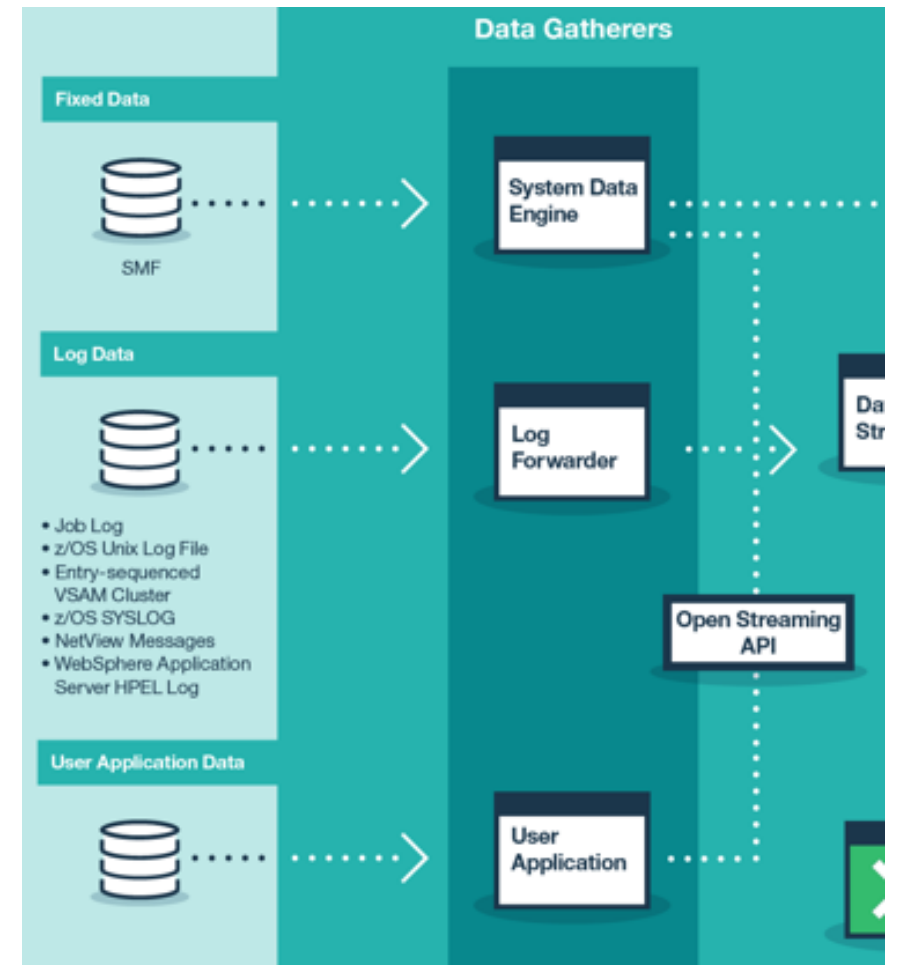
Data Gatherers

System Data Engine

- **Collect and process SMF data** (in near real-time or in batch)
- Supports multiple sources of SMF data (**SMF archive, Log Stream** or the **SMF in memory, user exit 83/84/85** for SYS1.MANx users)
 - **All** commonly used IBM SMF types are supported (**100+ SMF types**)
 - **CA TopSecret and ACF2** SMF records
 - **IMS Log** (Accounting, Security violation records, Transaction level statistics records and IMS Performance Analyzer Transaction Index records) and **IMS User Log**
- *Possibility to add **new data definition** through SDE language (*extensibility*)*
- SMF records can be converted into consumable data as **CSV or DB2 LOAD** format for easy ingestion by consumer

Log Forwarder

- Gathers a variety of log data and some VSAM file formats for Analytics Engines
 - CICS Transaction Server for z/OS logs,
 - NetView for z/OS messages.
 - WebSphere Application Server logs,
 - z/OS Syslog and USS SyslogD logs,
 - generic ZFS files and generic z/OS job output from a DD of a running job




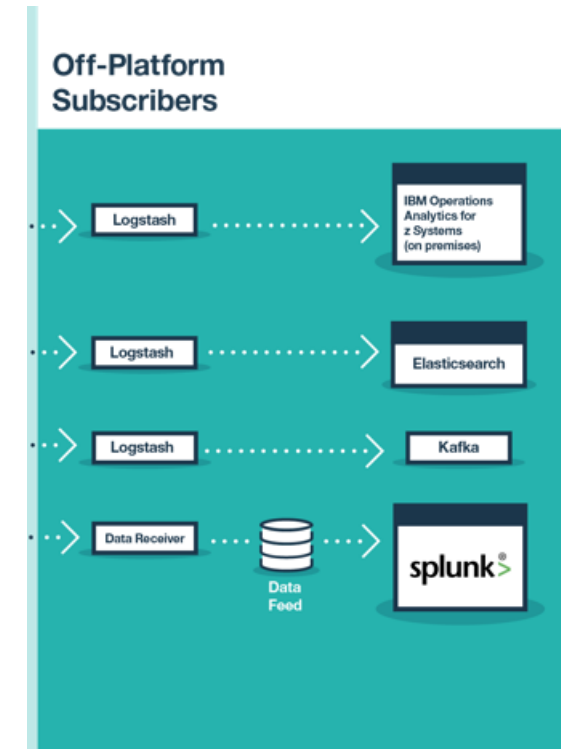
Open Streaming API (*extensibility*)

- Possibility to define streams for **user applications data** and to then push user application data down through those streams.
- Java and REXX API are available

Streaming Live Data

The Data Streamer controls the destination and format of data collected

- Receives data from the gatherers
- **Splits** the log data into individual messages for ingestion into analytic engines
- **Transforms** data messages into the right format for the destination platform (eg UTF-8 and other code pages)
- Transport is TCPIP with Optional **SSL**, HTTP or **HTTPS**.
- Data can be sent in **json wrapper** for ingestion by Logstash for storage and analysis
- Streams data both on and off platform
- **zIIP enabled** for cost savings (pure Java)



The Data Receiver writes any data it receives out into files

- Small Java program that can run on z/OS or on a distributed platform
- Data received are split in files by source type
- Enable easy ingestion from consumers like Splunk

IBM Common Data Provider for z System and Splunk

With IBM Common Data Provider for z Systems and Splunk clients can:

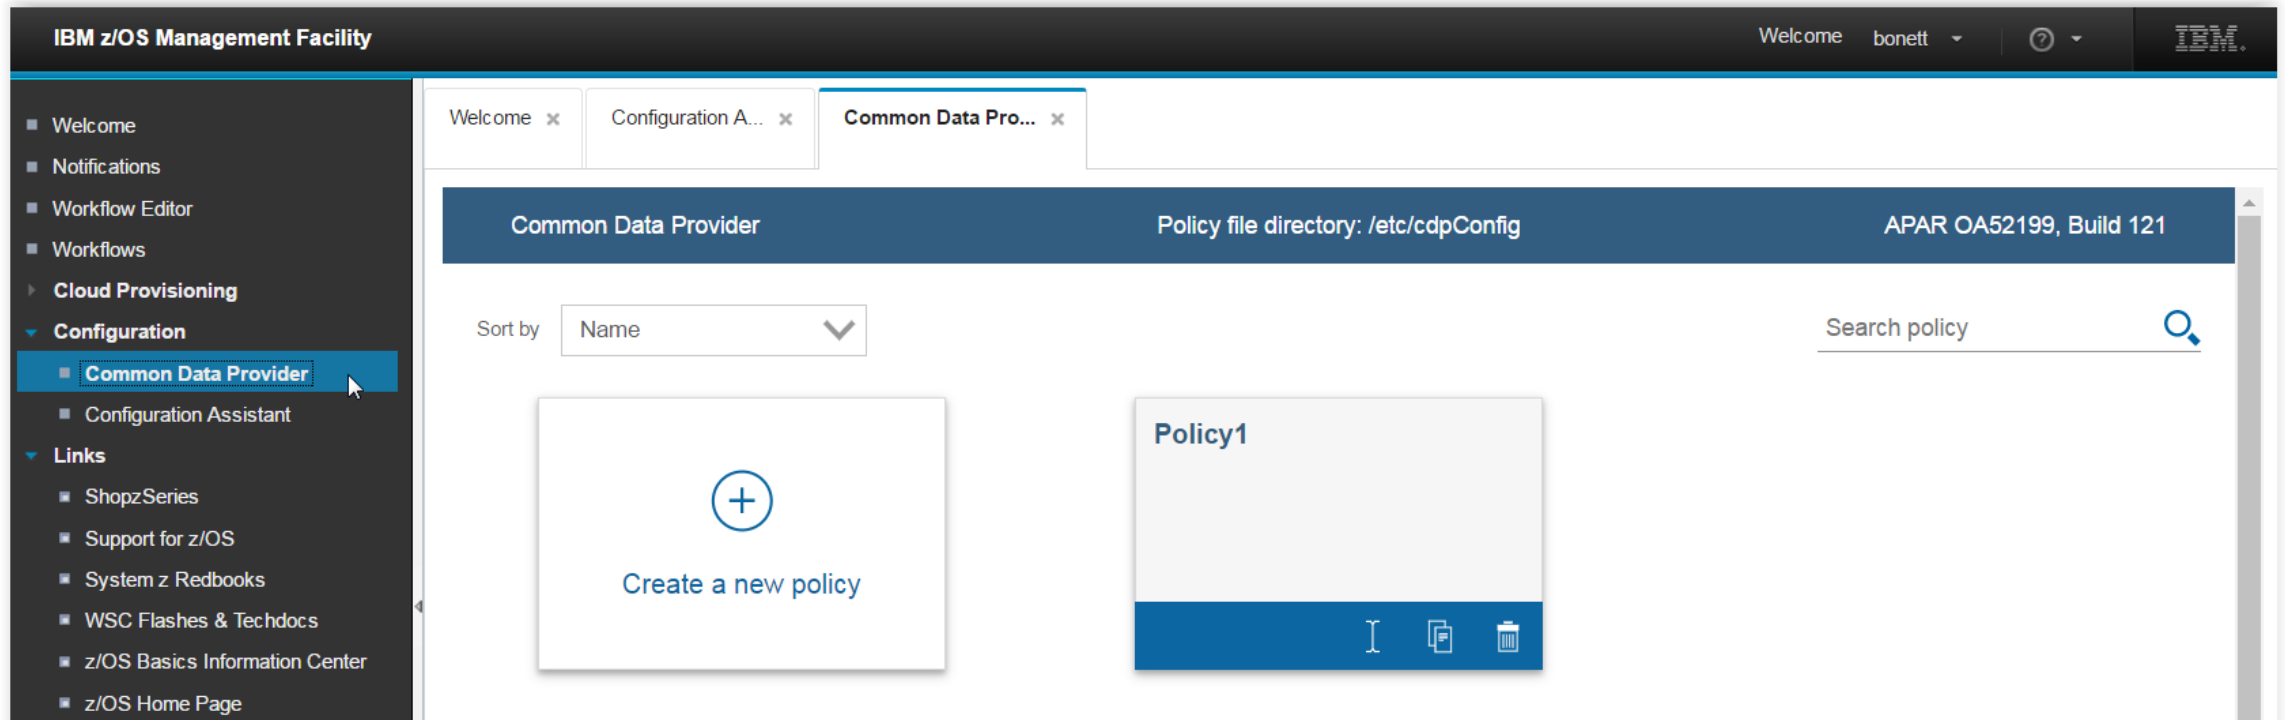
- **Gain insight** into **hybrid IT operations** by integrating IBM Z operational data with your public cloud and distributed operational data into a single analytics platform, **eliminating blind spots**
- **Visualize impacts** across your infrastructure from **continuously delivering** applications and application enhancements
- Stream the **widest range of SMF records and Z log data in “real-time”** to provide diagnostic and resource usage information.
- Build Custom **Enterprise Wide Splunk Apps** using Z Data.
- **Reduce** Splunk ingestion **costs** with **advanced filtering** and customer control of SMF and log data.
- **Save money** with the CDPz’s **fixed pricing model**



IBM Common Data Provider for z Systems

Product configuration

CDPz – Create a Policy



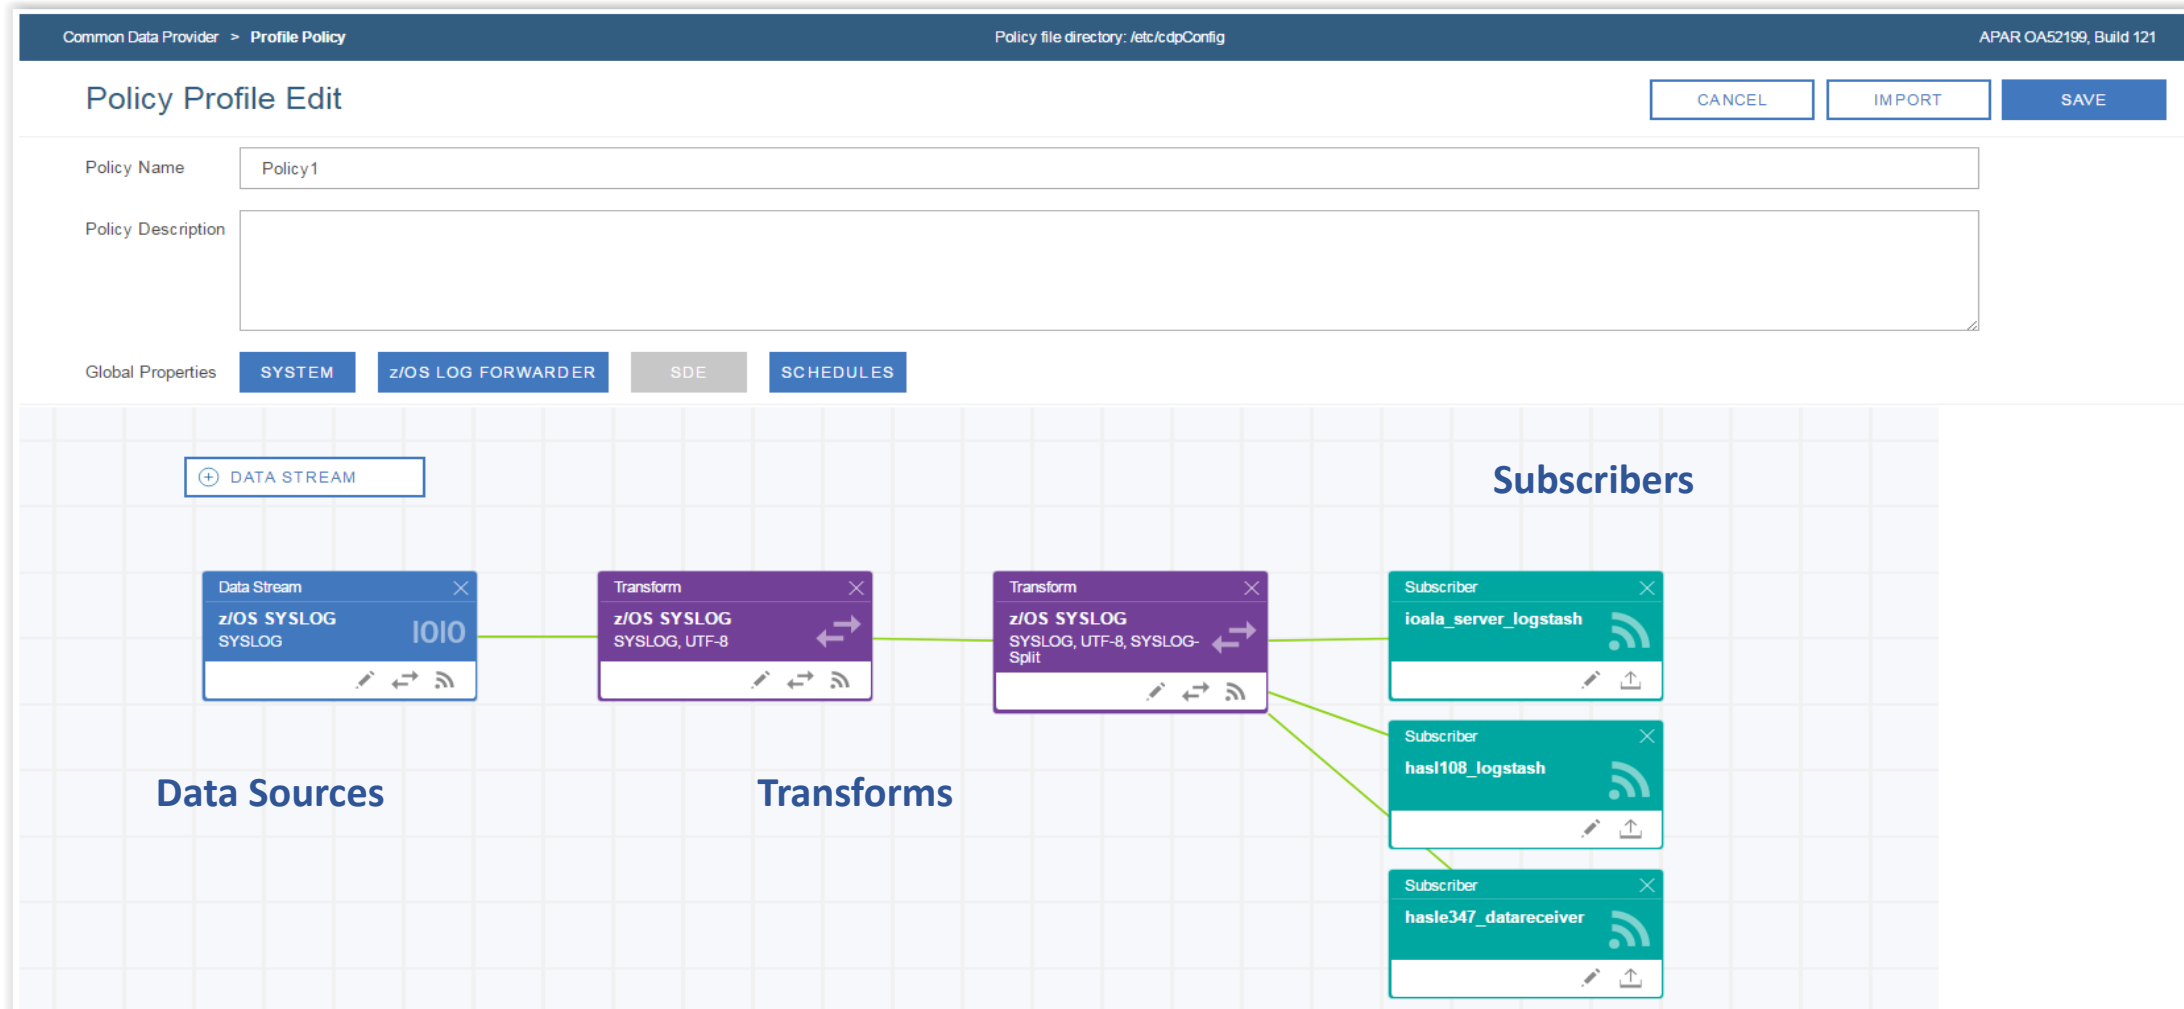
Create policies (sets of configurations files) for the Data Streamer task

Multiple polices can be defined, only one active at a time

Generates definitions stored in a set of files in configuration tool working directory

CDPz: Policy definition flow

Policy consists of data sources<->Transforms<->Subscribers (Consumers)



Policy definition flow: Data Sources

Policy consists of **data sources** <-> Transforms <-> Subscribers (Consumers)

- SMF Data

- Base Streams

- + z/OS
- + JES
- + SMF
- + RMF
- + Address Space
- + Device
- + SMS
- + Catalog
- + VSAM
- + Data Set
- + TSO



+ Security

+ DB2

- CICS

<input type="checkbox"/>	Name	Tags
<input type="checkbox"/>	SMF_110_0	CSV, Split
<input type="checkbox"/>	SMF_110_1	CSV, Split
<input type="checkbox"/>	SMF_110_1_FIELD	CSV, Split
<input type="checkbox"/>	SMF_110_1_DICT	CSV, Split
<input type="checkbox"/>	SMF_110_1_6	CSV, Split
<input type="checkbox"/>	SMF_110_E	CSV, Split
<input type="checkbox"/>	SMF_110_1_5	CSV, Split
<input type="checkbox"/>	SMF_110_2	CSV, Split
<input type="checkbox"/>	SMF_110_3	CSV, Split
<input type="checkbox"/>	SMF_110_4	CSV, Split
<input type="checkbox"/>	SMF_110_5	CSV, Split

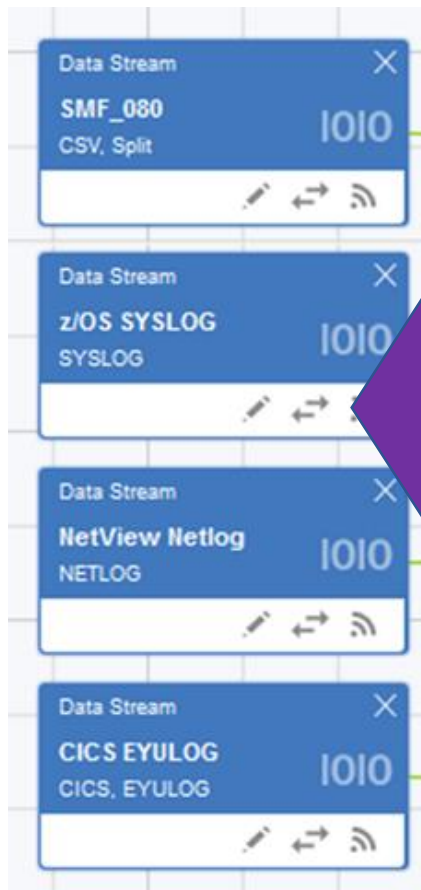


The screenshot shows a configuration panel for Data Streams. It contains four entries, each with a title, a description, and a value of 1010. Each entry has icons for edit, refresh, and share.

- Data Stream SMF_080**: CSV, Split
- Data Stream z/OS SYSLOG**: SYSLOG
- Data Stream NetView Netlog**: NETLOG
- Data Stream CICS EYULOG**: CICS, EYULOG

Policy definition flow: Transforms & Subscribers

Policy consists of data sources <-> **Transforms** <-> Subscribers (Consumers)



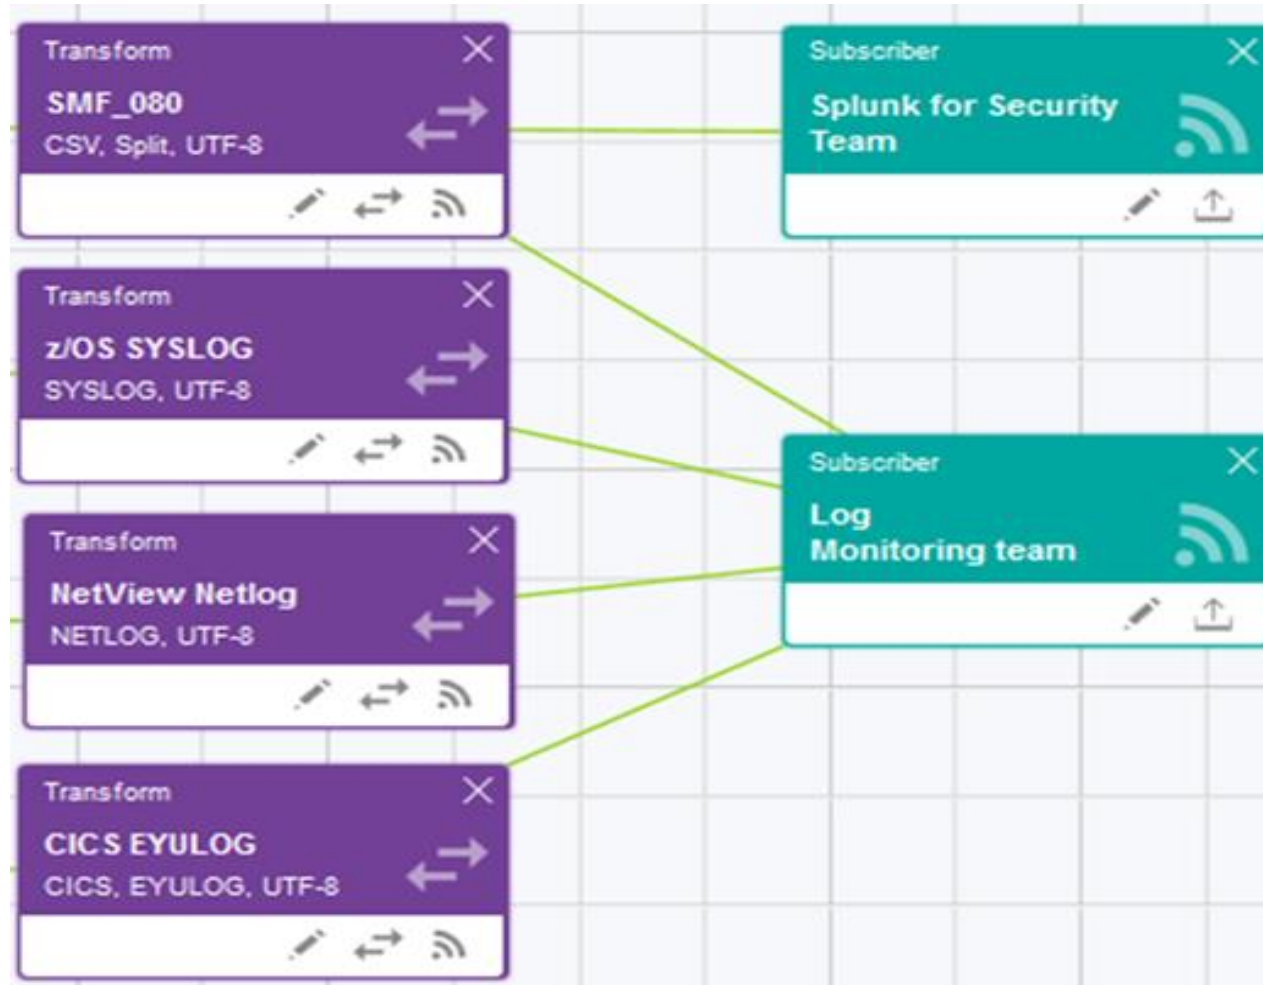
Transform data stream ✕

Transform Type	Description
<input type="radio"/> TRANSCRIBE	Transcribe the data to a different encoding
<input type="radio"/> CRLF Splitter	Split the data into multiple messages based on CRLF characters
<input type="radio"/> SYSLOG Splitter	Split SYSLOG data into individual messages
<input type="radio"/> SyslogD Splitter	Split SyslogD data into individual messages
<input type="radio"/> NetView Splitter	Split NetView data into individual messages
<input type="radio"/> EYULOG MDY Splitter	Split EYULOG data in MDY format into individual messages
<input type="radio"/> EYULOG DMY Splitter	Split EYULOG data in DMY format into individual messages
<input type="radio"/> EYULOG YMD Splitter	Split EYULOG data in YMD format into individual messages
<input type="radio"/> CICS MSGUSR MDY Splitter	Split CICS MSGUSR data in MDY format into individual messages
<input type="radio"/> CICS MSGUSR DMY Splitter	Split CICS MSGUSR data in DMY format into individual messages
<input type="radio"/> CICS MSGUSR YMD Splitter	Split CICS MSGUSR data in YMD format into individual messages
<input type="radio"/> WAS for zOS SYSOUT Splitter	Split WAS for zOS SYSOUT data into individual messages
<input type="radio"/> WAS for zOS SYSPRINT Splitter	Split WAS for zOS SYSPRINT data into individual messages
<input type="radio"/> WAS HPEL Splitter	Split WAS HPEL data into individual messages
<input type="radio"/> WAS SYSTEMOUT Splitter	Split WAS SYSTEMOUT data into individual messages
<input type="radio"/> FixedLength Splitter	Split records with a fixed record length into multiple messages
<input type="radio"/> Regex Filter	Filter messages from incoming streams based on a regex pattern
<input type="radio"/> Time Filter	Only allow packets sent within a particular schedule, and discard all others

Show all transforms
 TRANSFORM

Policy definition flow: Transforms & Subscribers

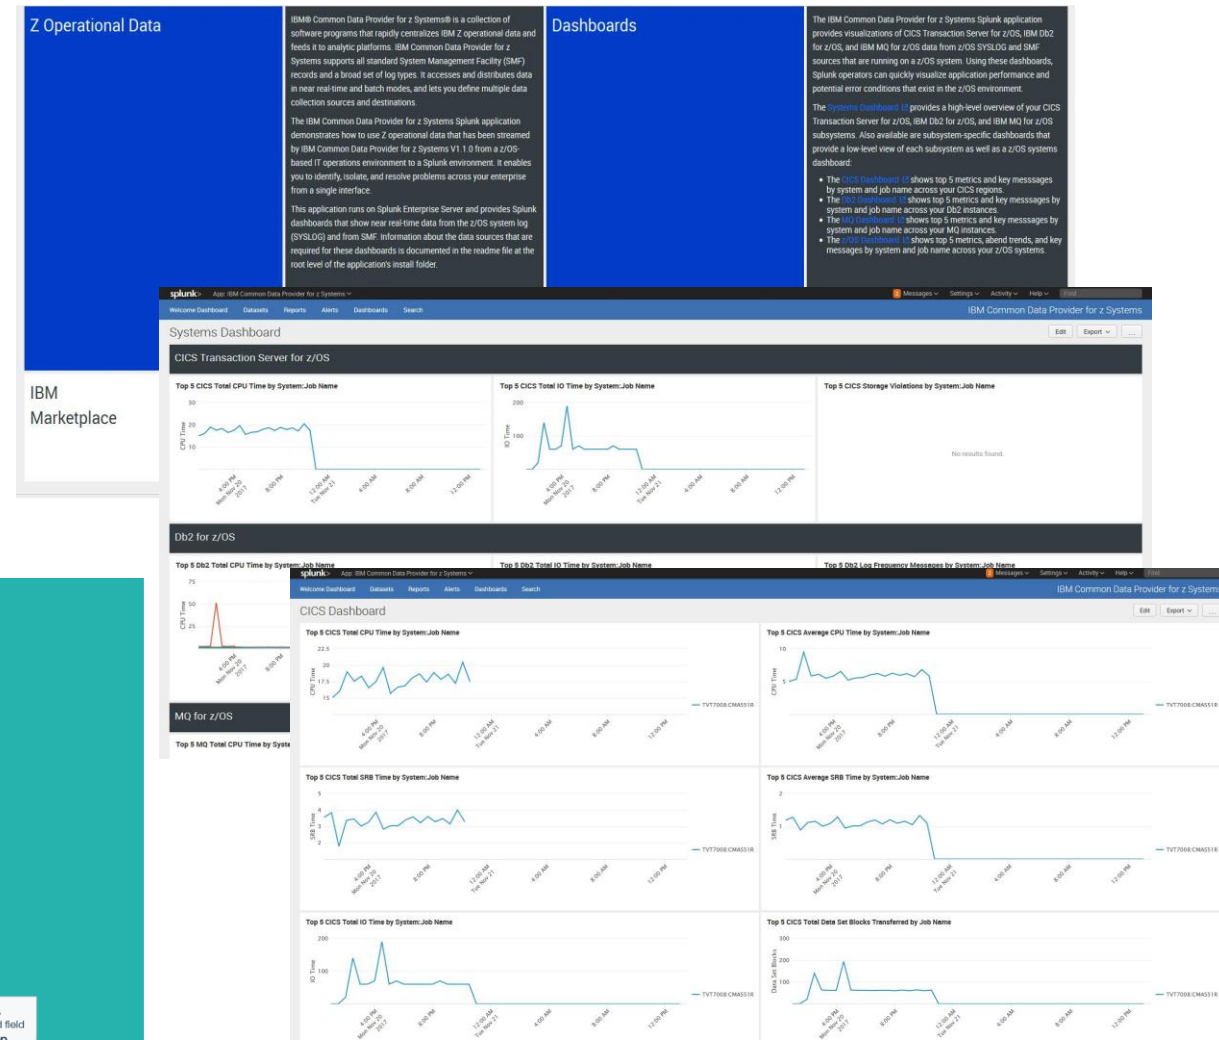
Policy consists of data sources<->Transforms<->**Subscribers** (Consumers)



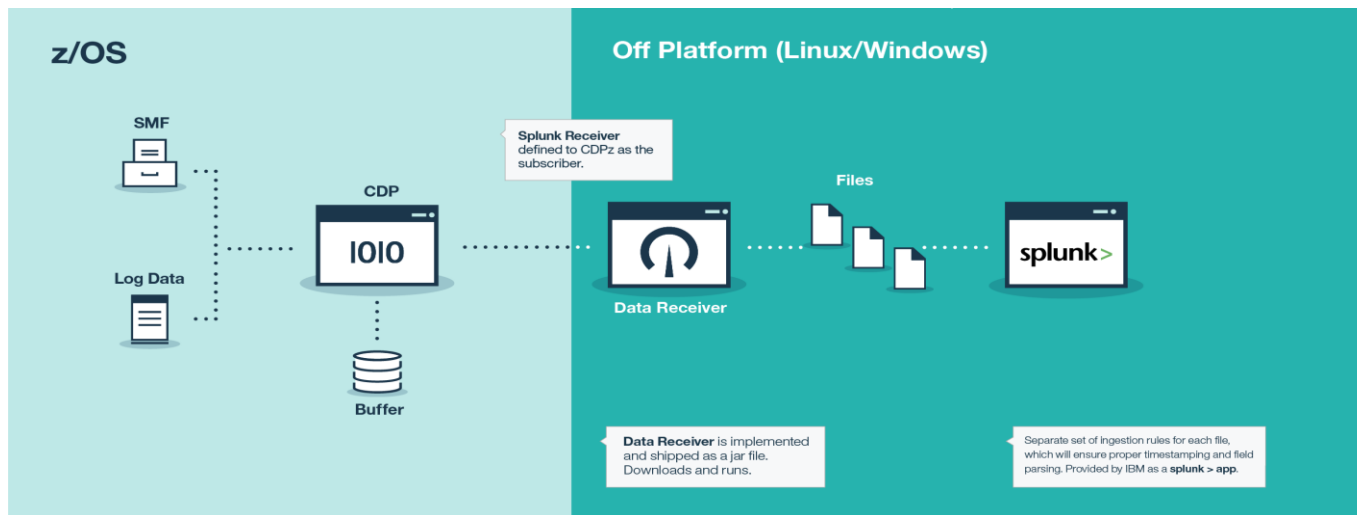
Common Data Provider stream data to Splunk

3 steps configuration

1. Download and install the CDPz Data Receiver
 - Set HOME and PATH environment variable, and Start it
2. Update your CDPz policy to add a subscriber pointing to the DataReceiver
3. Install the IBM CDPz Ingestion App for Splunk



The image shows two screenshots of Splunk dashboards. The top screenshot is titled 'Systems Dashboard' and displays three line charts: 'Top 5 CICS Total CPU Time by System:Job Name', 'Top 5 CICS Total IO Time by System:Job Name', and 'Top 5 CICS Storage Violations by System:Job Name'. The bottom screenshot is titled 'CICS Dashboard' and displays six line charts: 'Top 5 CICS Total CPU Time by System:Job Name', 'Top 5 CICS Average CPU Time by System:Job Name', 'Top 5 CICS Total SRB Time by System:Job Name', 'Top 5 CICS Average SRB Time by System:Job Name', 'Top 5 CICS Total IO Time by System:Job Name', and 'Top 5 CICS Total Data Set Blocks Transferred by Job Name'. The dashboards are part of the 'IBM Common Data Provider for z/OS' application.

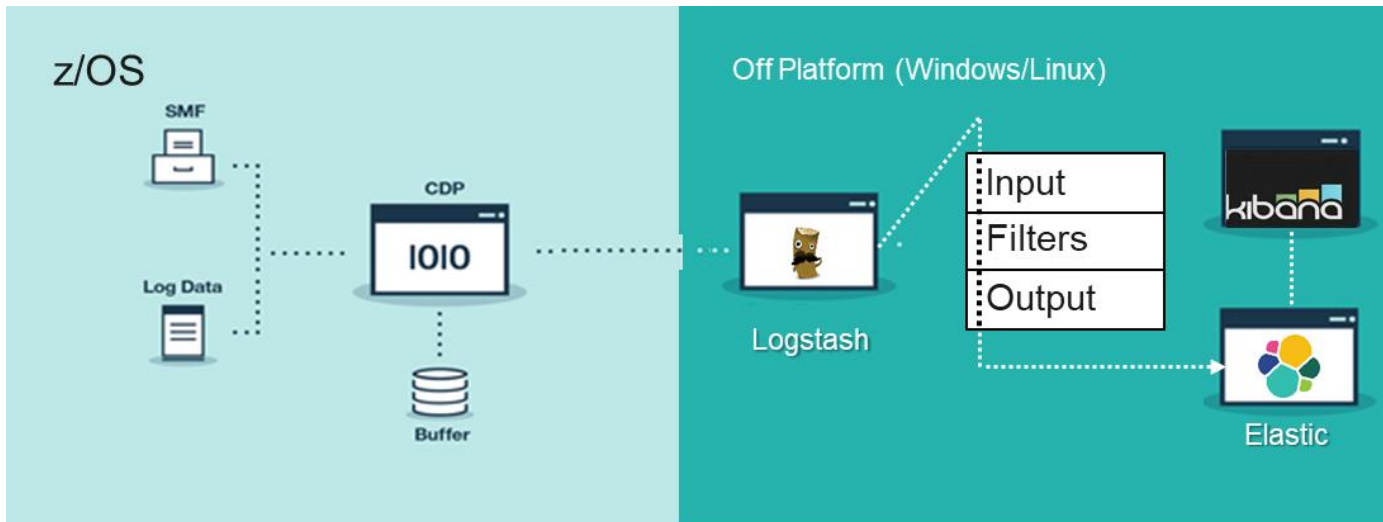
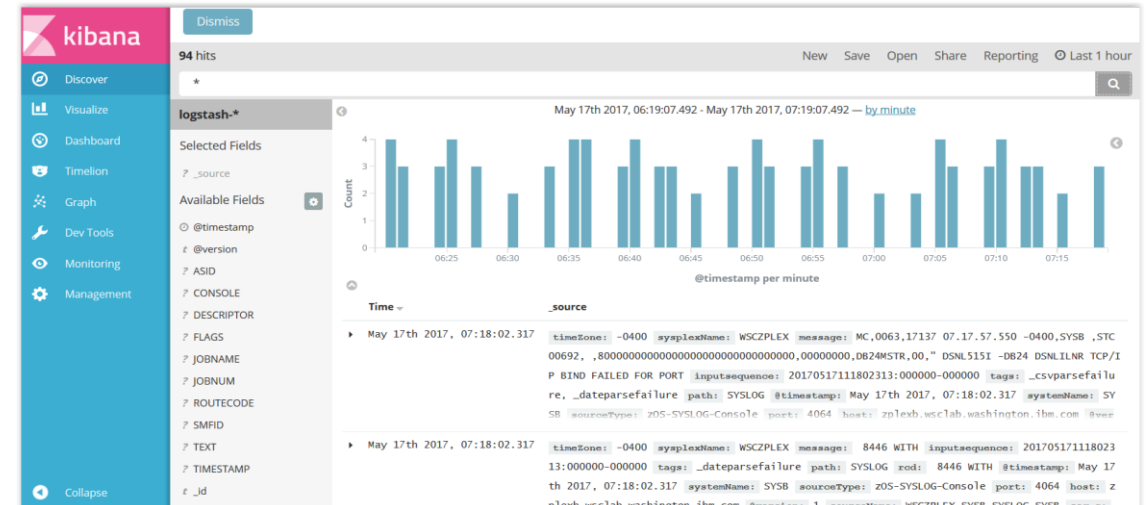


CDPz sample dashboards for Splunk are available on [Splunkbase](https://splunkbase.com)

Common Data Provider stream data to ELK

3 step configuration:

1. Download and expand the ELK ingestion kit
2. Copy the config files to the Logstash directory
 - Logstash must be using a configuration directory
 - Copy Input, Output and Filter files as appropriate
3. Restart Logstash



Common Data Provider supplies Logstash config files for:

1. Input – TCP listener
2. Filters – For each data type, define the CSV format
3. Elastic output - index => sourceType, Sysplex, Date

CDPz Sample Dashboards on ELK are available in [Mainframe Dev](#)

SDE Language: New record definition

DEFINE RECORD statement to define the SMF record type that describes the layout of the new record

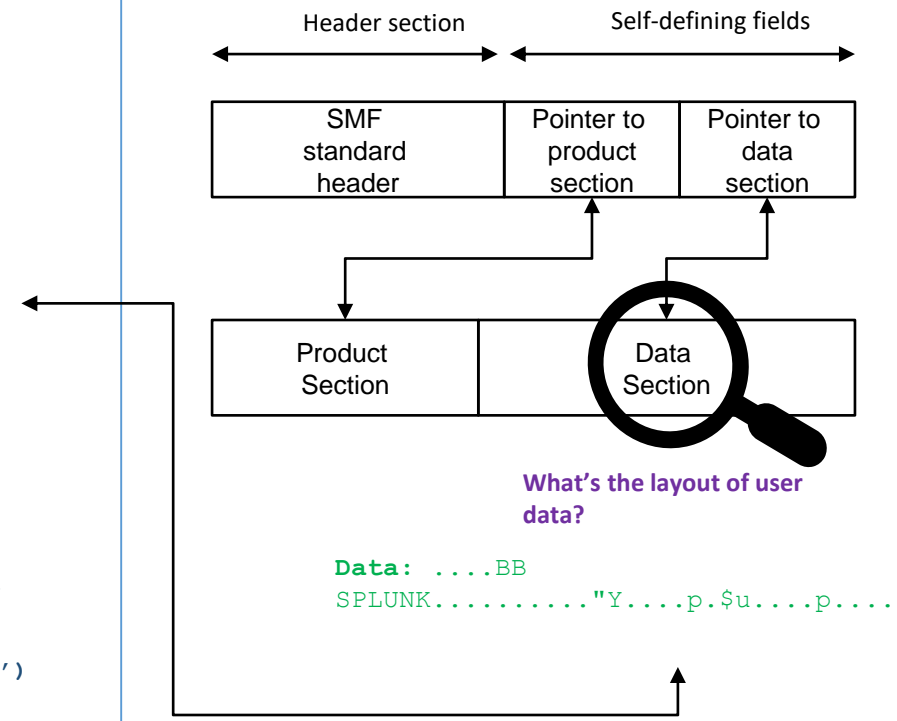
```

DEFINE RECORD SMF_IMS_C2C2
  VERSION 'CDP.110'
  IN LOG SMF
  IDENTIFIED BY SMIMSRTY = &IMS_SMF_RECTYPE
                 AND SMIMSSTY = &IMS_SMF_RECSTYP
                 AND SMIMSSRC = &IMS_SMF_SRCID
                 AND SMIMSID1 = 7

-----
-- Start of SMF header
-----
FIELDS (
  SMIMSLen LENGTH 2 BINARY,
  .....
-----
-- Start of CDP product section
-----
SECTION PRODUCT
  .....
-----
-- Start of IMS user record
-----
SECTION SMF_IMS_C2C2
  .....
  FIELDS (
    RECLL    LENGTH 2 BINARY,  -- Length of log record
    RECZZ    LENGTH 2 HEX,     -- QSAM reserved bits
    RECTYPE  LENGTH 1 BINARY,  -- Record type (X'C2')
    RECSUB   LENGTH 1 BINARY,  -- Record subtype (x'C2')
    RECSPL   LENGTH 4 CHAR,    -- Filler
    RECSPL   LENGTH 6 CHAR,    -- C' SPLUNK'
    .....
  )

```

This record definition shows an example of an IMS user log record "C2C2" definition



Record Level Filtering: What record to collect?

DEFINE UPDATE statement to specify:

- how to process a specific record type
- how to store the result in the target of the update

```

DEFINE UPDATE SMF_101_1_FT
VERSION 'CDP.110'
FROM SMF_101_1
WHERE (SUBSTR(QWHCEUTX,1,2) = 'MG')
      OR (QWHCAID = 'U@MUPJ2')
TO &IBM_UPDATE_TARGET
&IBM_CORRELATION
AS &IBM_FILE_FORMAT SET(ALL);
    
```

This update definition performs record level filtering on the SMF type 101 subtype 1 record (DB2)

SMF records read from Exit, Log stream, or In-memory resource

SM1LEN	SM101SGD	...	QWHCAID	...	QWHCEUTX	...	QPKGPKG	QPKGPKNF	QPKGPKNL	Record collected?
901	0		U@MUPJ2		muecps.bin		4	1	4	✓
901	0		U#MGXP01		MG005		6	1	6	✓
901	0		U@FCPJ1		TF008		25	21	25	✗
901	0		U@PYPJ4		ib85pr1d		8	10	20	✗
901	0		U@MUPJ2		MG121		10	1	10	✓

Record Level Filtering: where condition

DEFINE UPDATE statement with WHERE condition enables the development of **complex data streams with calculations** (like the CICS KPI data stream delivered in CDPz)

An example: CICS long running or abend

```

DEFINE UPDATE SMF_110_1_KPI_SB
VERSION 'CDP.110'
FROM SMF_CICS_T
WHERE (SUBSTR(TRAN,1,2) = 'QF' AND INTERVAL(START, STOP) > 2.0)
      OR (ABCODEO <> ' ')
      OR (ABCODEC <> ' ')
  
```

The WHERE condition selects CICS records for transactions meeting ONE of the following criteria:

- a. Any “QF” transactions whose **response time is greater than 2.0 seconds**
- b. Transaction ended with **abend**

Field level filtering: What's the desired output?

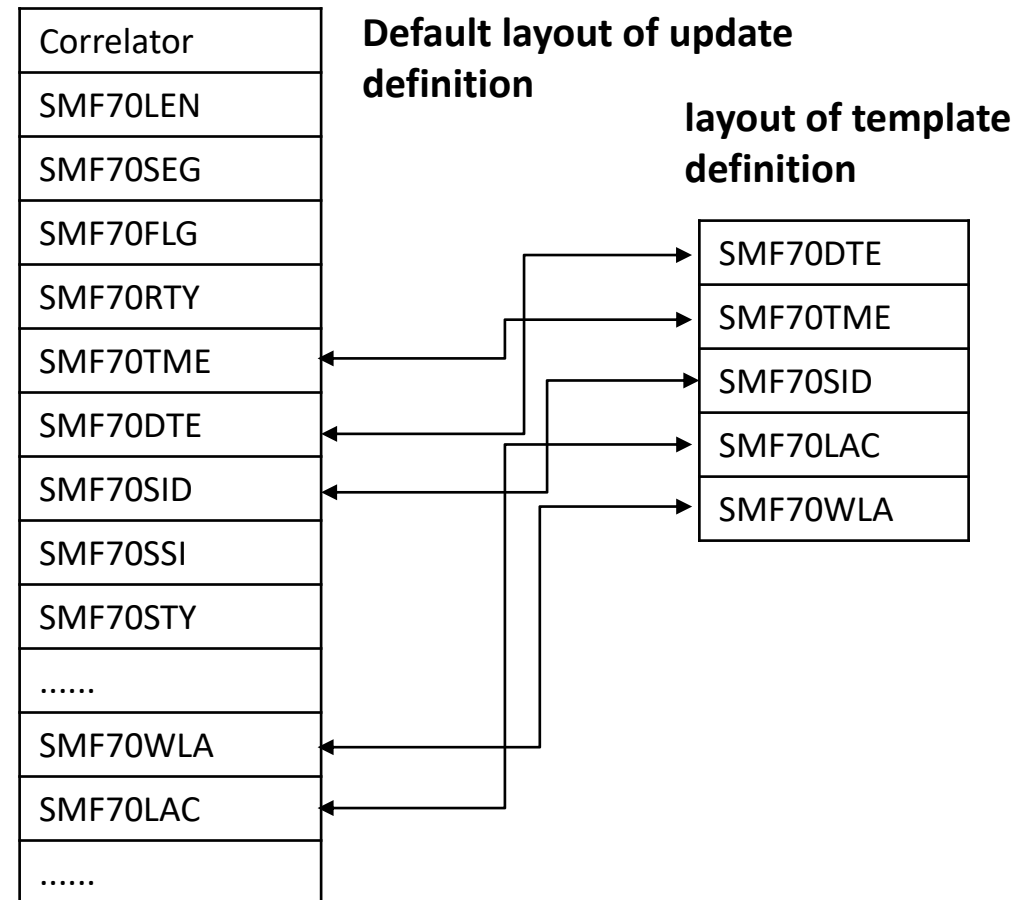
DEFINE TEMPLATE statement can be used in conjunction with an update object to do any of the following:

- Create additional output files with the same data, or a subset of the data.
- Change the order of the fields in the output file.
- Alter the formatting of the fields in the output file.

```
DEFINE TEMPLATE SMF_070 FOR  
SMF_070  
  ORDER (  
    SMF70DTE,  
    SMF70TME,  
    SMF70SID,  
    SMF70LAC,  
    SMF70WLA  
  )  
  TO &IBM_UPDATE_TARGET  
  &IBM_CORRELATION  
  AS &IBM_FILE_FORMAT;
```



This template definition only sends 5 fields from the original record to the subscriber



IBM Common Data Provider for z Systems

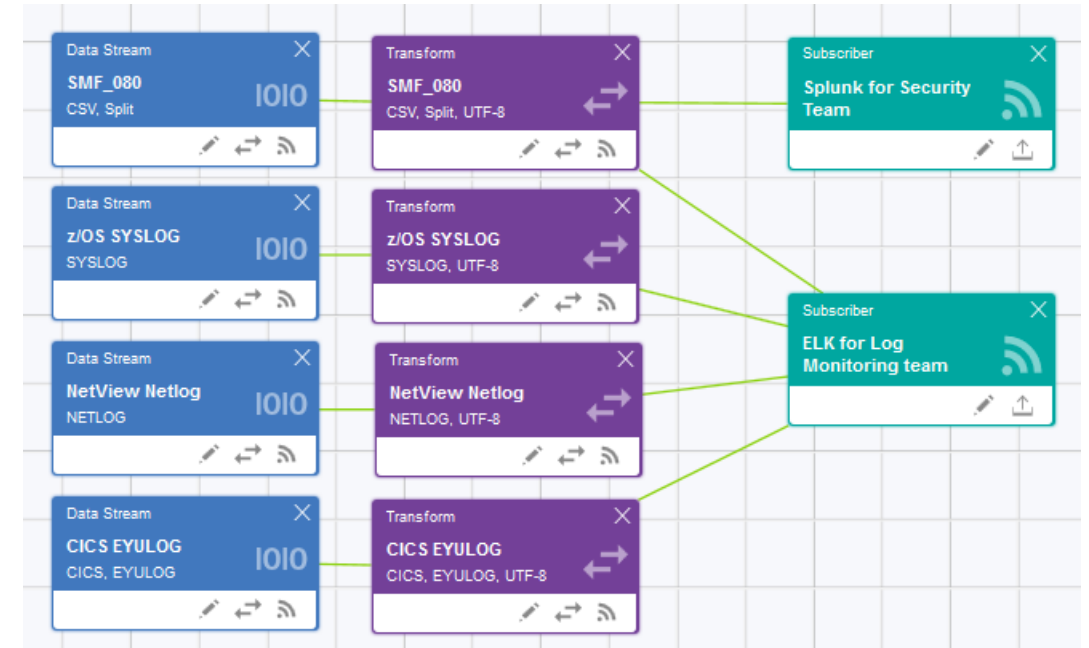
User scenarios

Send data to different targets: a European retail scenario

Customer background and needs:

- **Splunk** is already used in the Enterprise for Network, Device and Open Systems Security
- Need to **integrate z Systems** security data into Splunk for Enterprise security team
- Mainframe IT Operation team plan to use **ELK** for real time log monitoring

Customer looking for **one tool able** to satisfy the needs



IBM Common Data Provider for z System is the tool able to read the data once and send to different targets

Different filtering/transformation options can be applied to the data before sending the same data to different targets.

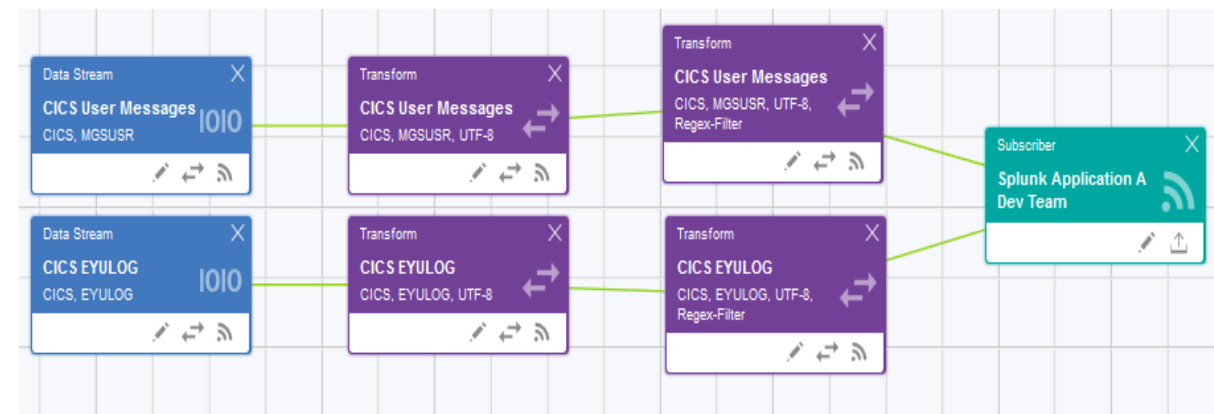
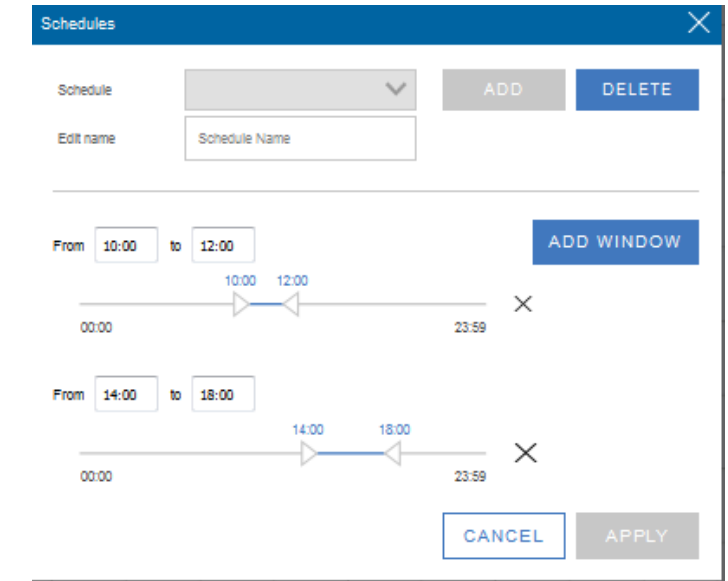
Enable a devOps scenario: a US Health Insurance Company

Customer background and needs:

- *Application A* is a customer mobile application with components on open systems as well as leveraging CICS on z Systems
- **Splunk** is already used by *Application A* development team to analyze *Application A* open systems application logs
- *Application A* development team **need** to have mainframe CICS log related to their own application in Splunk to make **end2end problem determination**
- *Application A* test of new features are performed in specific time windows

IBM Common Data Provider for z System is the solution being able to:

- Gather CICS log data only in specific time windows
- Filter CICS log data based on the content (like a custom filter referring to *Application A*)
- Select Send only selected data to the Splunk Server accessible to *Application A* dev Team



Rabobank use case



Instant Payments Infrastructure leveraging IBM Z data and Splunk IT Service Intelligence. Process payments in 5 seconds 24/7/365.

Visualize and predict disruptions to the business before customers are impacted

Reduce mean time to repair saving time in war room situations

IBM Common Data Provider for z Systems

Cross-product integrations

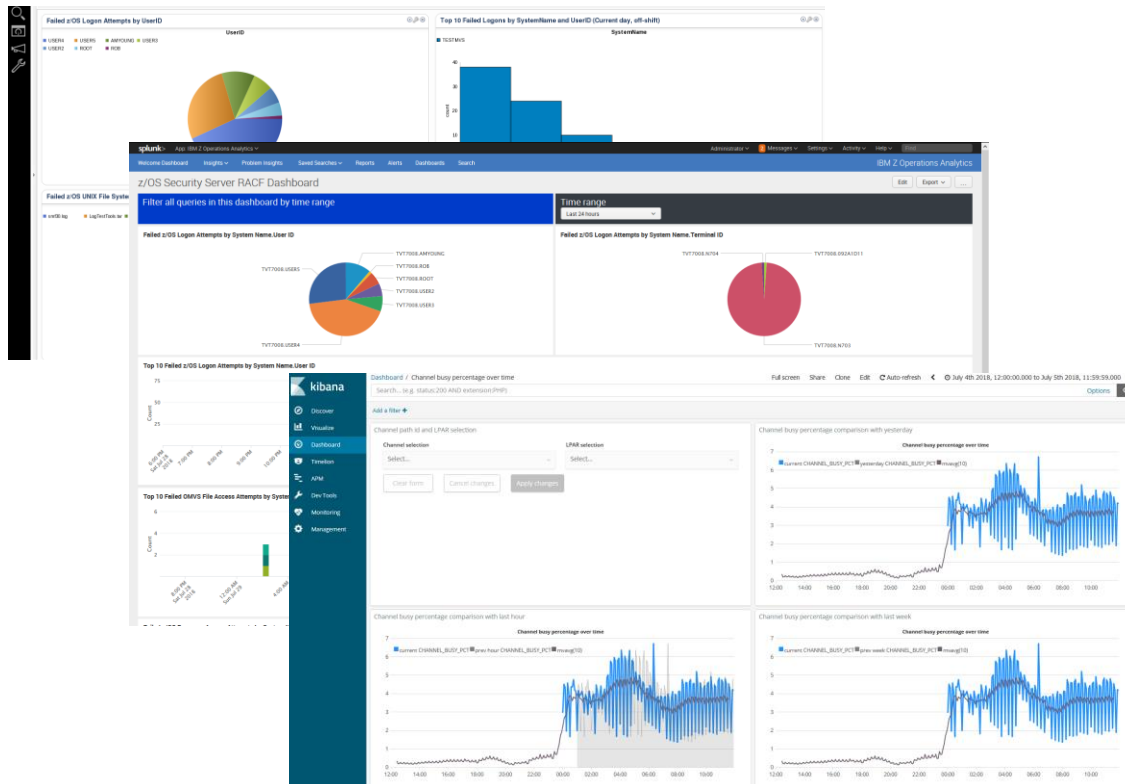
IBM Z Operations Analytics includes CDPz

Dashboards and Searches – Insight into IBM Z Operational Data

Use out-of-the-box dashboards on IBM, Splunk and Elastic or build your own with a click of the mouse

Domain-specific ‘Quick Searches’ available out-of-the-box

- Based on the combined experience of **subject matter experts, support teams and customers.**
- **Immediate value** out of the box.
- Easy to modify or create and **save your own.**



Db2 for z/OS

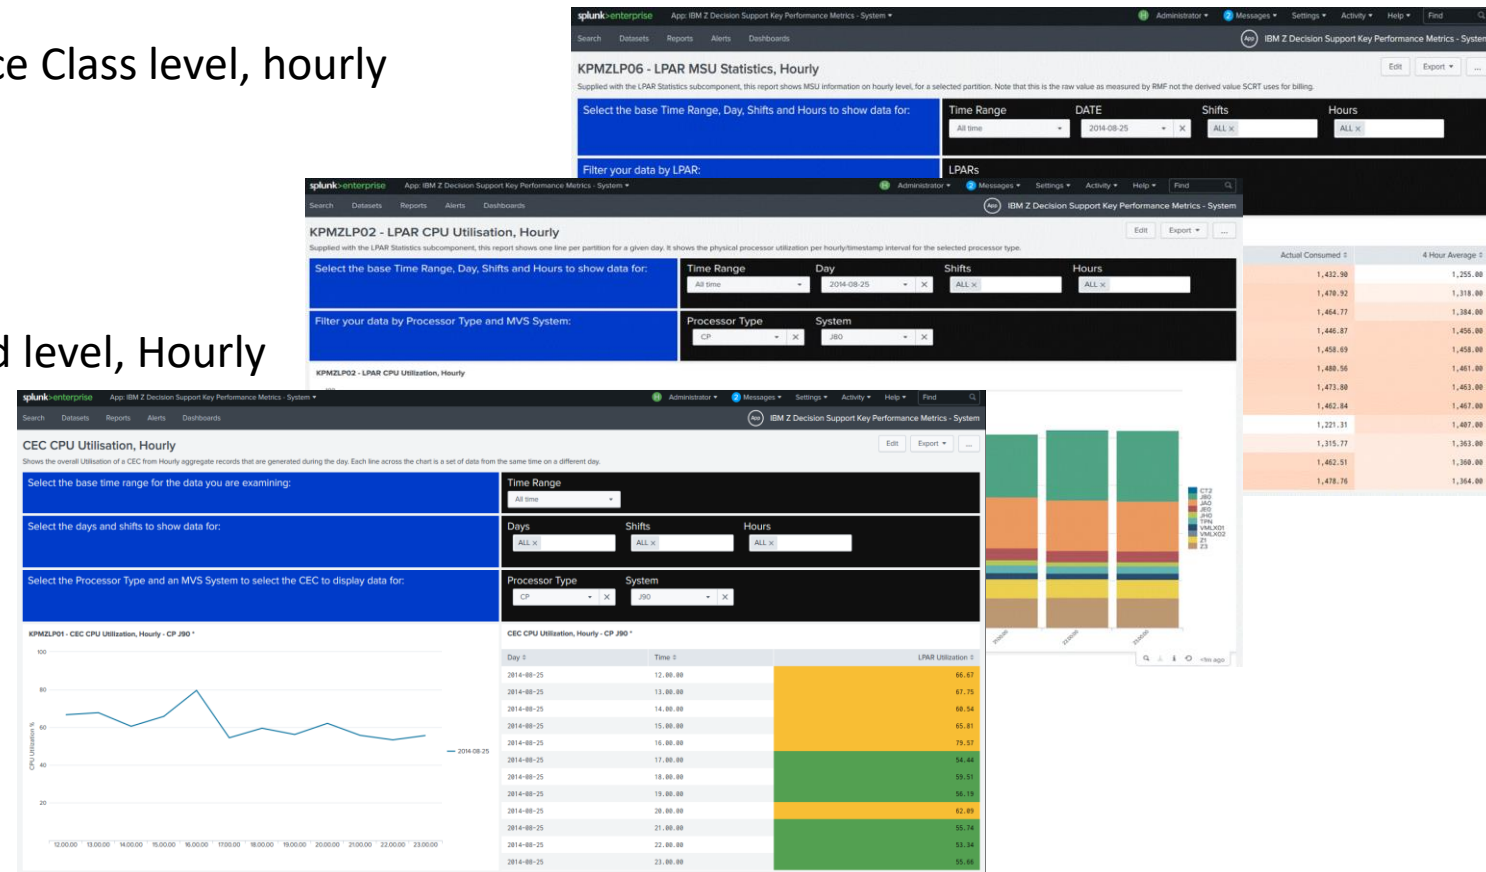
- Db2 Messages
- Db2 Action, Decision, or Error Messages
- Db2 Data Set Messages
- Db2 Data Sharing Messages
- Db2 Lock Conflict Messages
- Db2 Log Data Set Messages
- Db2 Log Frequency Messages
- Db2 Pool Shortage Messages
- Db2 Job Performance Metrics

IBM Z Decision Support integration with CDPz

Gain Z Operational data visibility in Splunk leveraging the integration of IBM Z Decision Support with IBM Common Data Provider for z systems.

Dashboards on

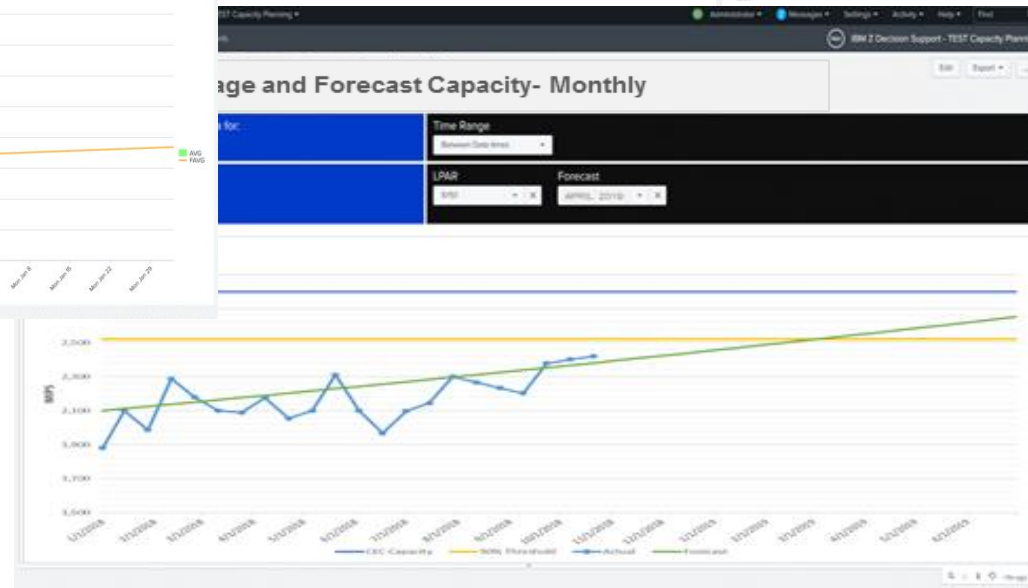
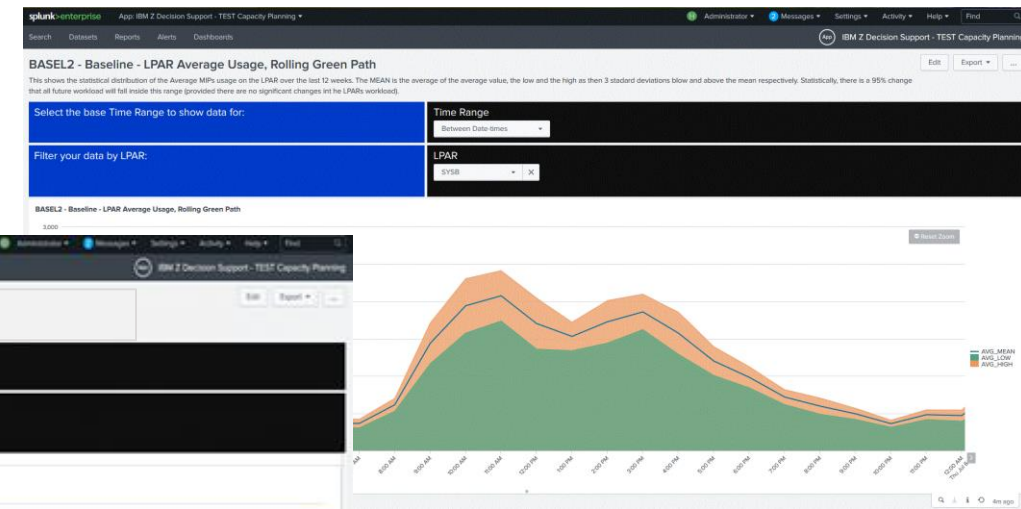
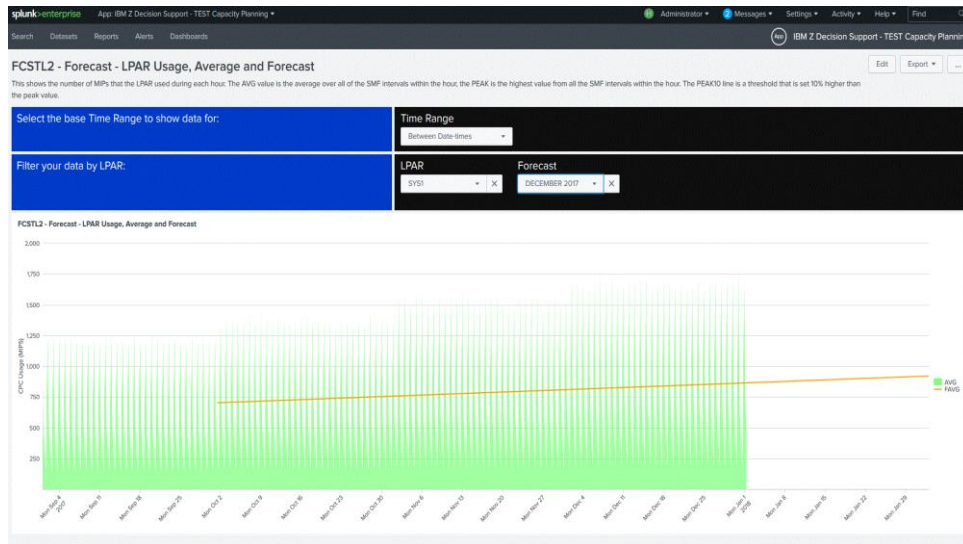
- **CPU Utilization @ CEC, LPAR, Workload, Service Class level, hourly**
- **LPAR Share, Hourly**
- **LPAR CPU Logical Utilization, Hourly**
- **LPAR 4 Hour MSU Utilization, Hourly**
- **LPAR MSU Statistics, Hourly**
- **System Storage Summary @ System, Workload level, Hourly**
- **Job CPU reports**
- **CICS Transaction Performance, Hourly**



Session 0L: Conquer performance challenges with IBM Z Decision Support insight and analytics Wednesday 12:00 PM

IBM Z Decision Support for Capacity Planning integration with CDPz

Gain Forecasting and Predictive Analysis visibility in Splunk leveraging the integration of IBM Z Decision Support for Capacity Planning with IBM Common Data Provider for z systems.

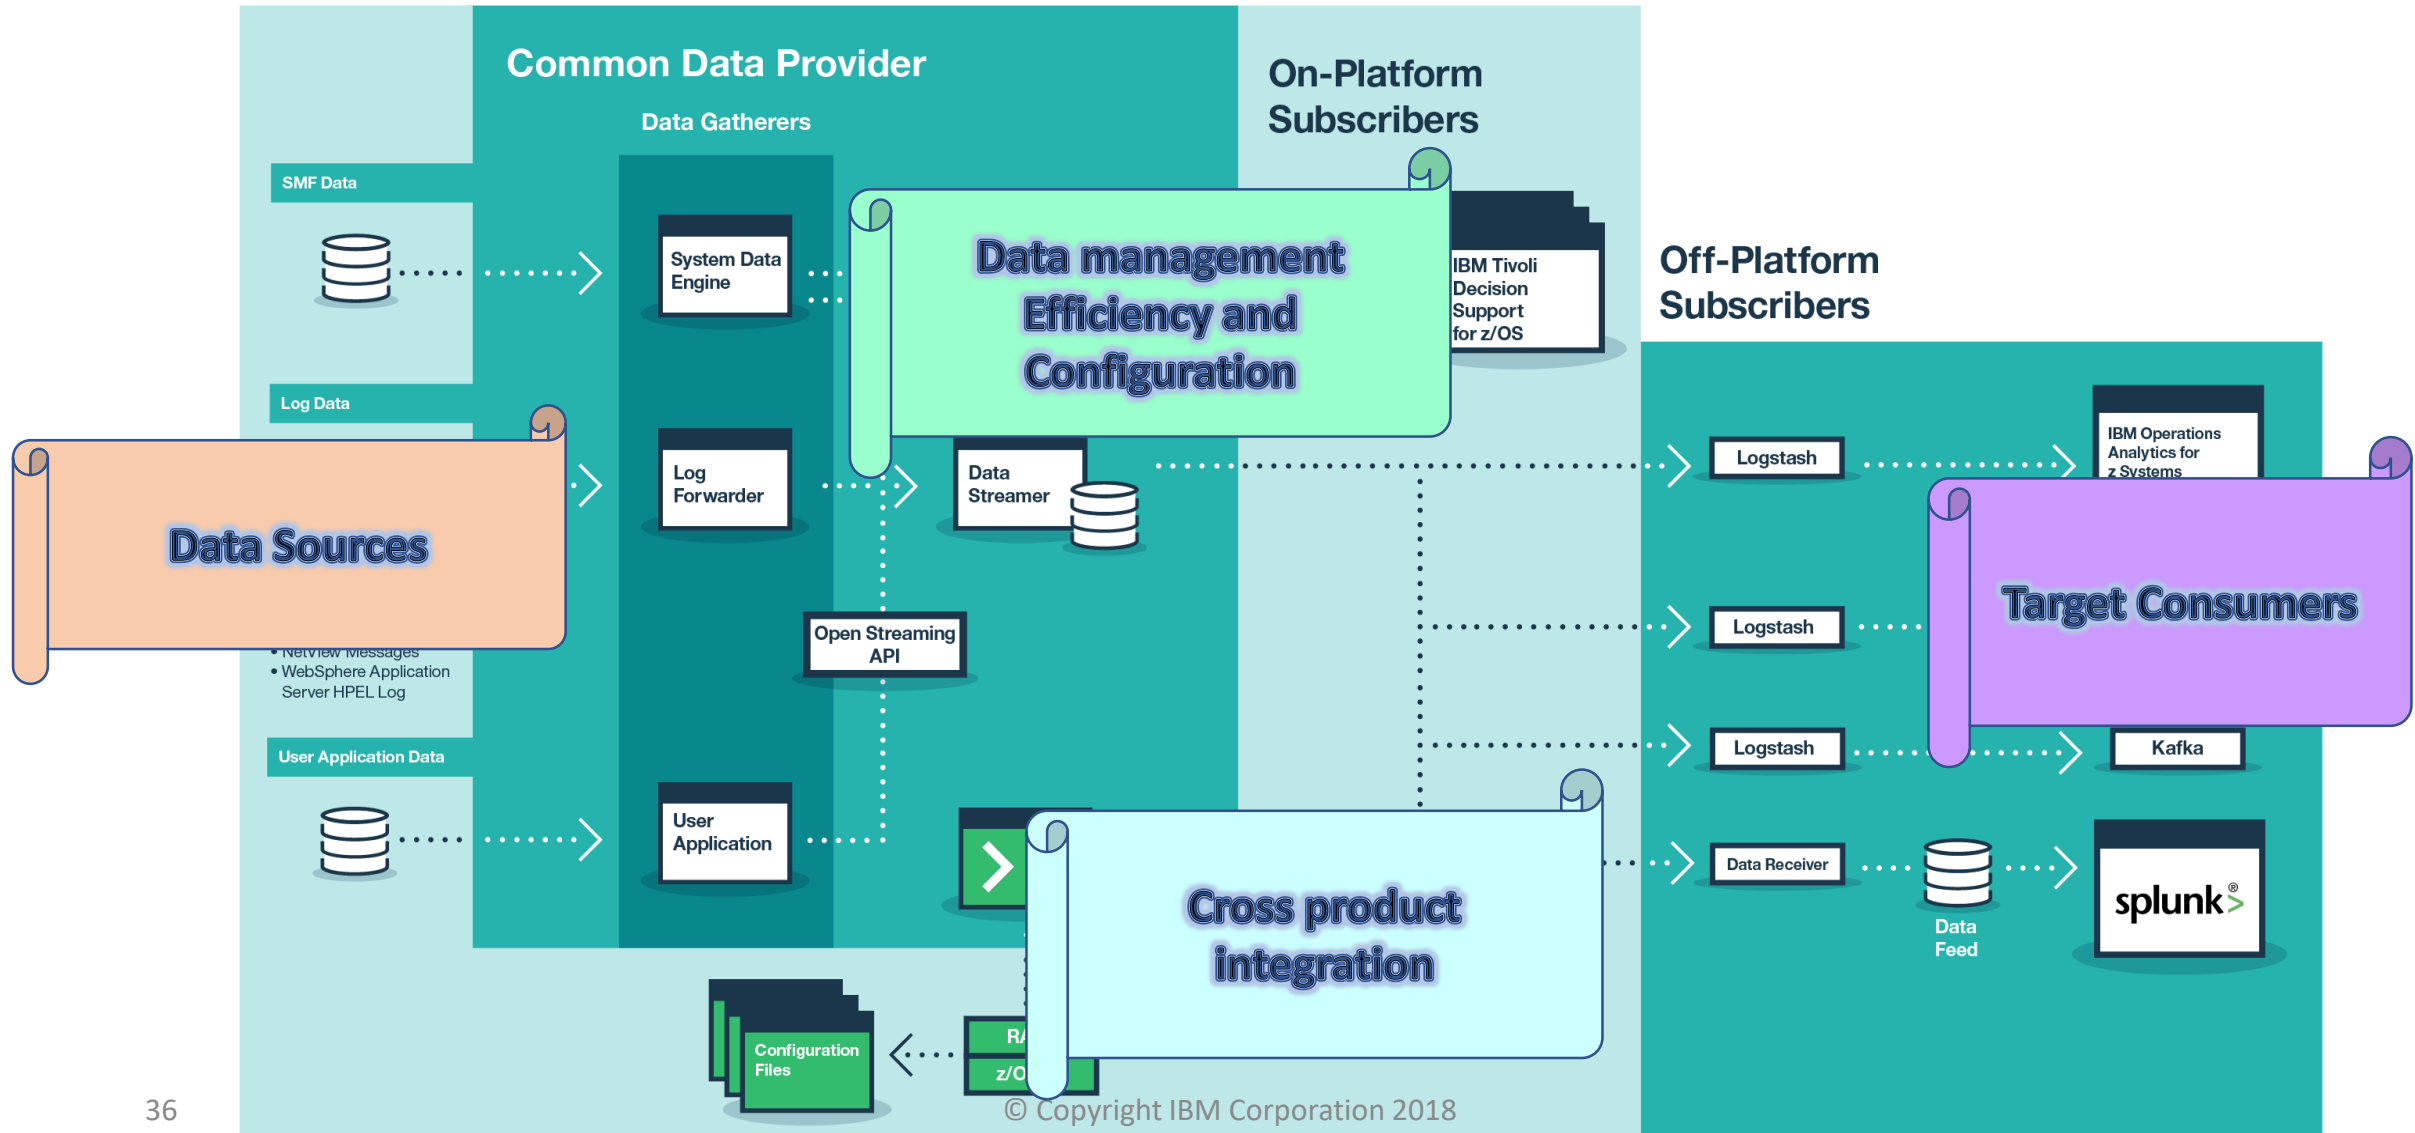


Session OM: Conquer performance challenges with IBM Z Decision Support insight and analytics Wednesday 14:00 PM

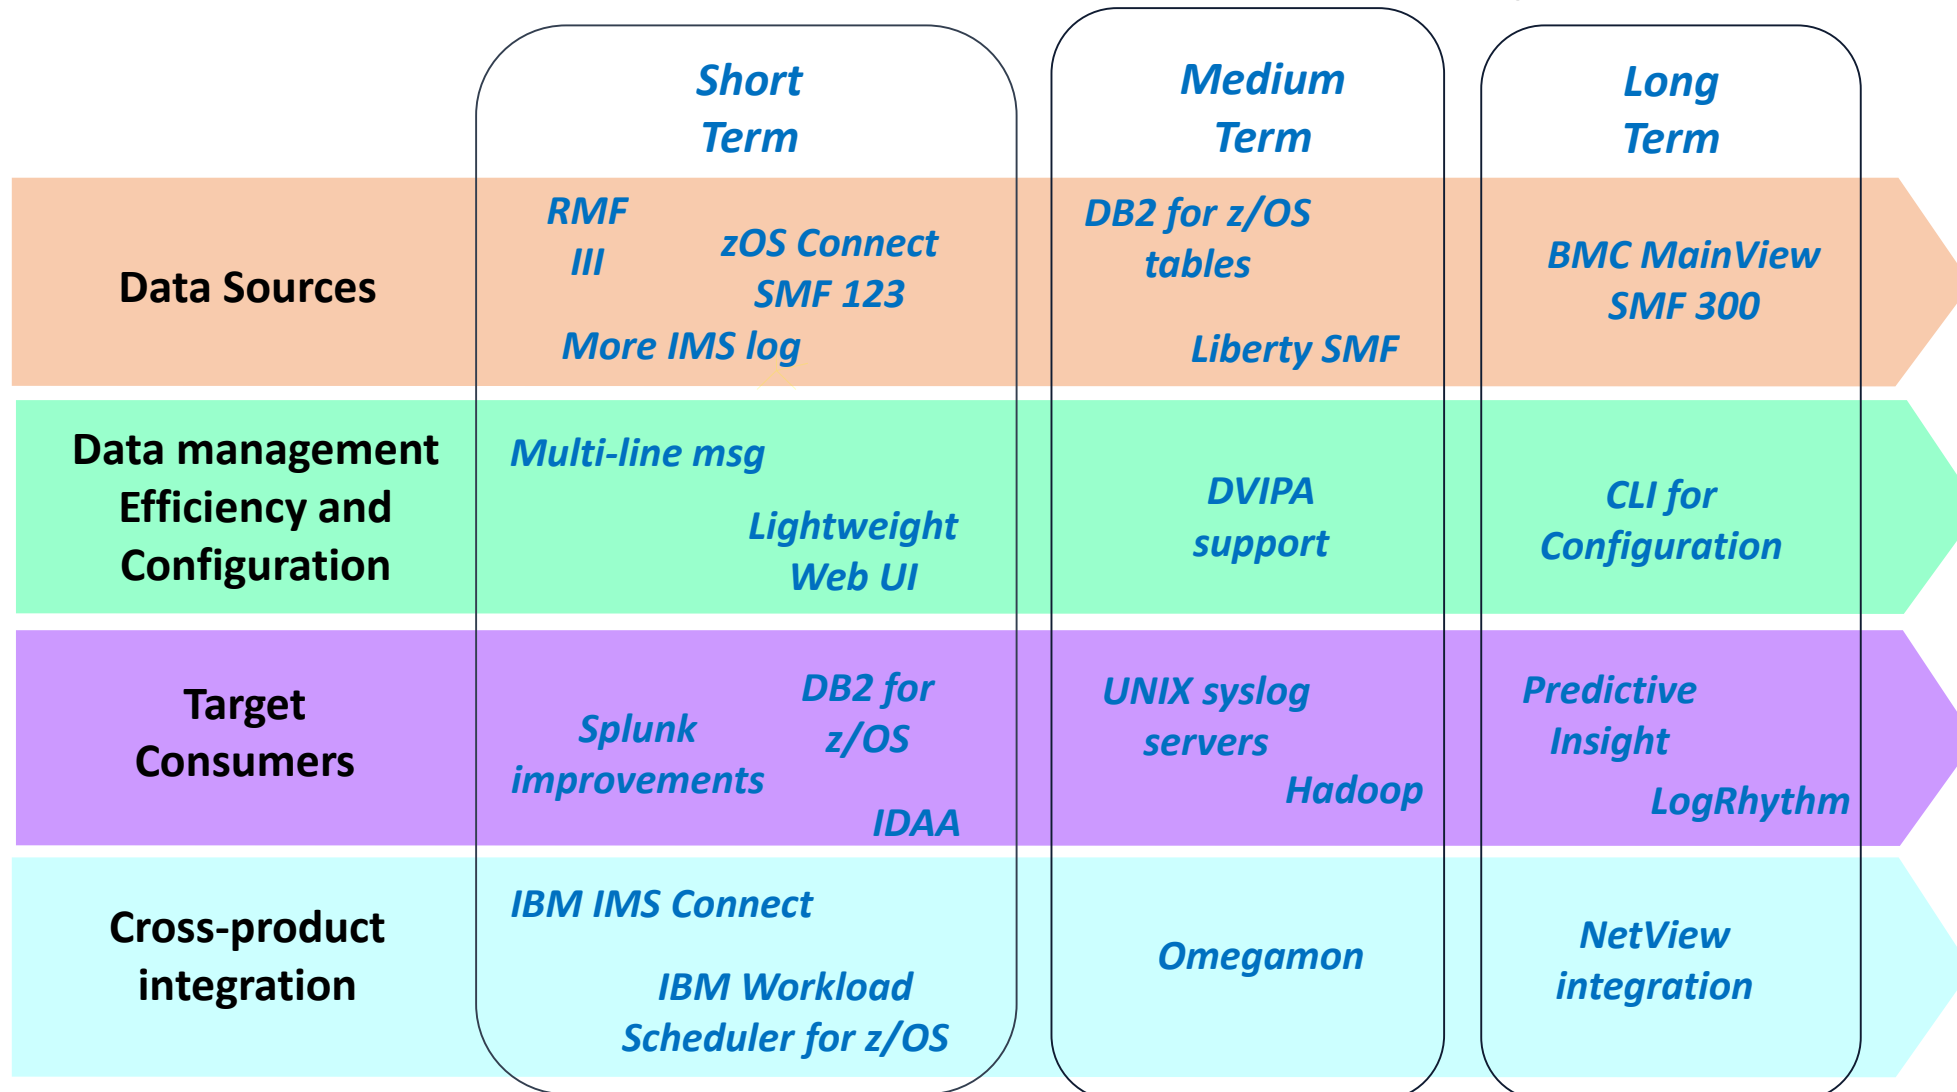
IBM Common Data Provider for z Systems

Roadmap & Strategy

Common Data Provider Strategy



Common Data Provider Roadmap



Further information

IBM Common Data Provider for Z Systems marketplace site <http://ibm.biz/CDPzInfo>

IBM Common Data Provider 1.1 Knowledge Center <https://ibm.biz/CDPzDoc>
Contains links to program directories (SMP/E installation) and Users Guide (product setup, customization, and integrations)

IBM Common Data Provider on DeveloperWorks <http://ibm.biz/CDPzWiki>
FAQs, documentations, service information

White Paper on integrating CDPz with Splunk
<https://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102713>

CDPz Sample Dashboard for Splunk on SplunkBase <http://ibm.biz/CDPzSamples>

White Paper on Integrating CDPz with ELK
<https://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102722>

CDP Sample Dashboard on Elastic Stack
<https://developer.ibm.com/mainframe/2018/03/08/ibm-common-data-provider-ibm-z-insight-dashboards-elastic-stack/>

We want your feedback!

- Please submit your feedback online at
 - <http://conferences.gse.org.uk/2018/feedback/ok>



- Paper feedback forms are also available from the Chair person

- This session is OK

