

# Unlocking Performance Monitoring Data for Analytics via Improved Observability Prepared for SPARTA

Jim Porell

[jporell@rocketsoftware.com](mailto:jporell@rocketsoftware.com)

June 2023



Rocket has rebranded – update spartanc.org 😊

---



# Prep Questions for SPARTA

Are you using either Docker or OpenShift zCX containers yet?

Does anyone in your organization use an Analytics monitoring tool: Elastic, Splunk, Instana, Datadog, Dynatrace, etc?

Do you attend the OMEGAMON Tech Summit virtual meetings? (Sep 2021, May 2022, Oct 2022)

- Next one is June 27<sup>th</sup> <https://community.ibm.com/community/user/ibmz-and-linuxone/blogs/katie-higgs2/2023/05/25/join-us-at-the-virtual-omegamon-technical-summit-o>

Do you leverage the IBM Community for z AIOps?

<https://community.ibm.com/community/user/ibmz-and-linuxone/groups/topic-home?CommunityKey=da213ad6-12c7-4f30-8c92-5ed2b2e1249e>



# Executive Summary

---

Customers have already proven to **benefit financially** and **reduce their time to resolve issues** through AI and ML

AI/ML is not a product...it's technology that can be applied to any product

OMEGAMON is embracing AI/ML to improve its handling of performance management

Customer input is needed to improve the training of models. **IBM and Rocket are looking for sponsor users for this new journey**

**I want you to enjoy this...nothing earth shattering except where the new technology can take us!**



# Agenda

---

ODP re-visited – a demo

- Elastic
- Grafana
- Mobile

Instana

- Youtube demo

OMEGAMON AI Insights Demo

# Challenges across the Organization



Application Developer

*"I have to keep up with application changes as we expand our markets worldwide and adopt more users."*




LOB Manager

*"I need to get my business results fast and accurate. What's going on ?"*



IT Manager

*"Performance problems seem to appear without warning and deep technical skills are hard to find."*



*"It is challenging to address performance issues with application workloads that always seem to be changing."*



CICS Sysprog

*"I'm getting too many performance alerts. Need to be able to pinpoint the ones that are most important."*




QA Manager


*"We can't test for changing workloads in today's digital economy since we don't have enough resources."*

# IBM Z AIOps will focus on three primary use cases


3 pillars to enable IBM Z clients to adopt enterprise observability



**Cross-portfolio integrations**  
to reduce the time to resolve  
operational issues



**Integrate with IBM Automation**  
to enhance hybrid cloud Observability  
and Proactive Incident Management



**Evolve AIOps**  
by embedding analytics and  
AI across the solutions


Integrated workflows to streamline the  
capabilities across the zAIOps portfolio

Observe and optimize capabilities  
across IBM Z and hybrid cloud

New AIOps use cases to improve time  
to resolution and address skill gaps

## Observability framework

IBM Z




- **Detect** Proactively identify IT Operations issues
- **Decide** Rapidly perform root cause analysis with AI
- **Act** Collaboratively resolve critical IT issues through automation

# OMEGAMON Data Provider Overview

# Observability leads to Insights driving Decisions

These insights can be used to drive decision making with regards to improving system performance as well as driving predicting analytics to solve potential problems before they fully manifest themselves.





# OMEGAMON Data Provider Overview

There is a wealth of subsystem data available within and across the OMEGAMON family today



With businesses looking to leverage analytics, artificial intelligence and machine learning to automate processes and act on the myriad of operational data available within their enterprises.

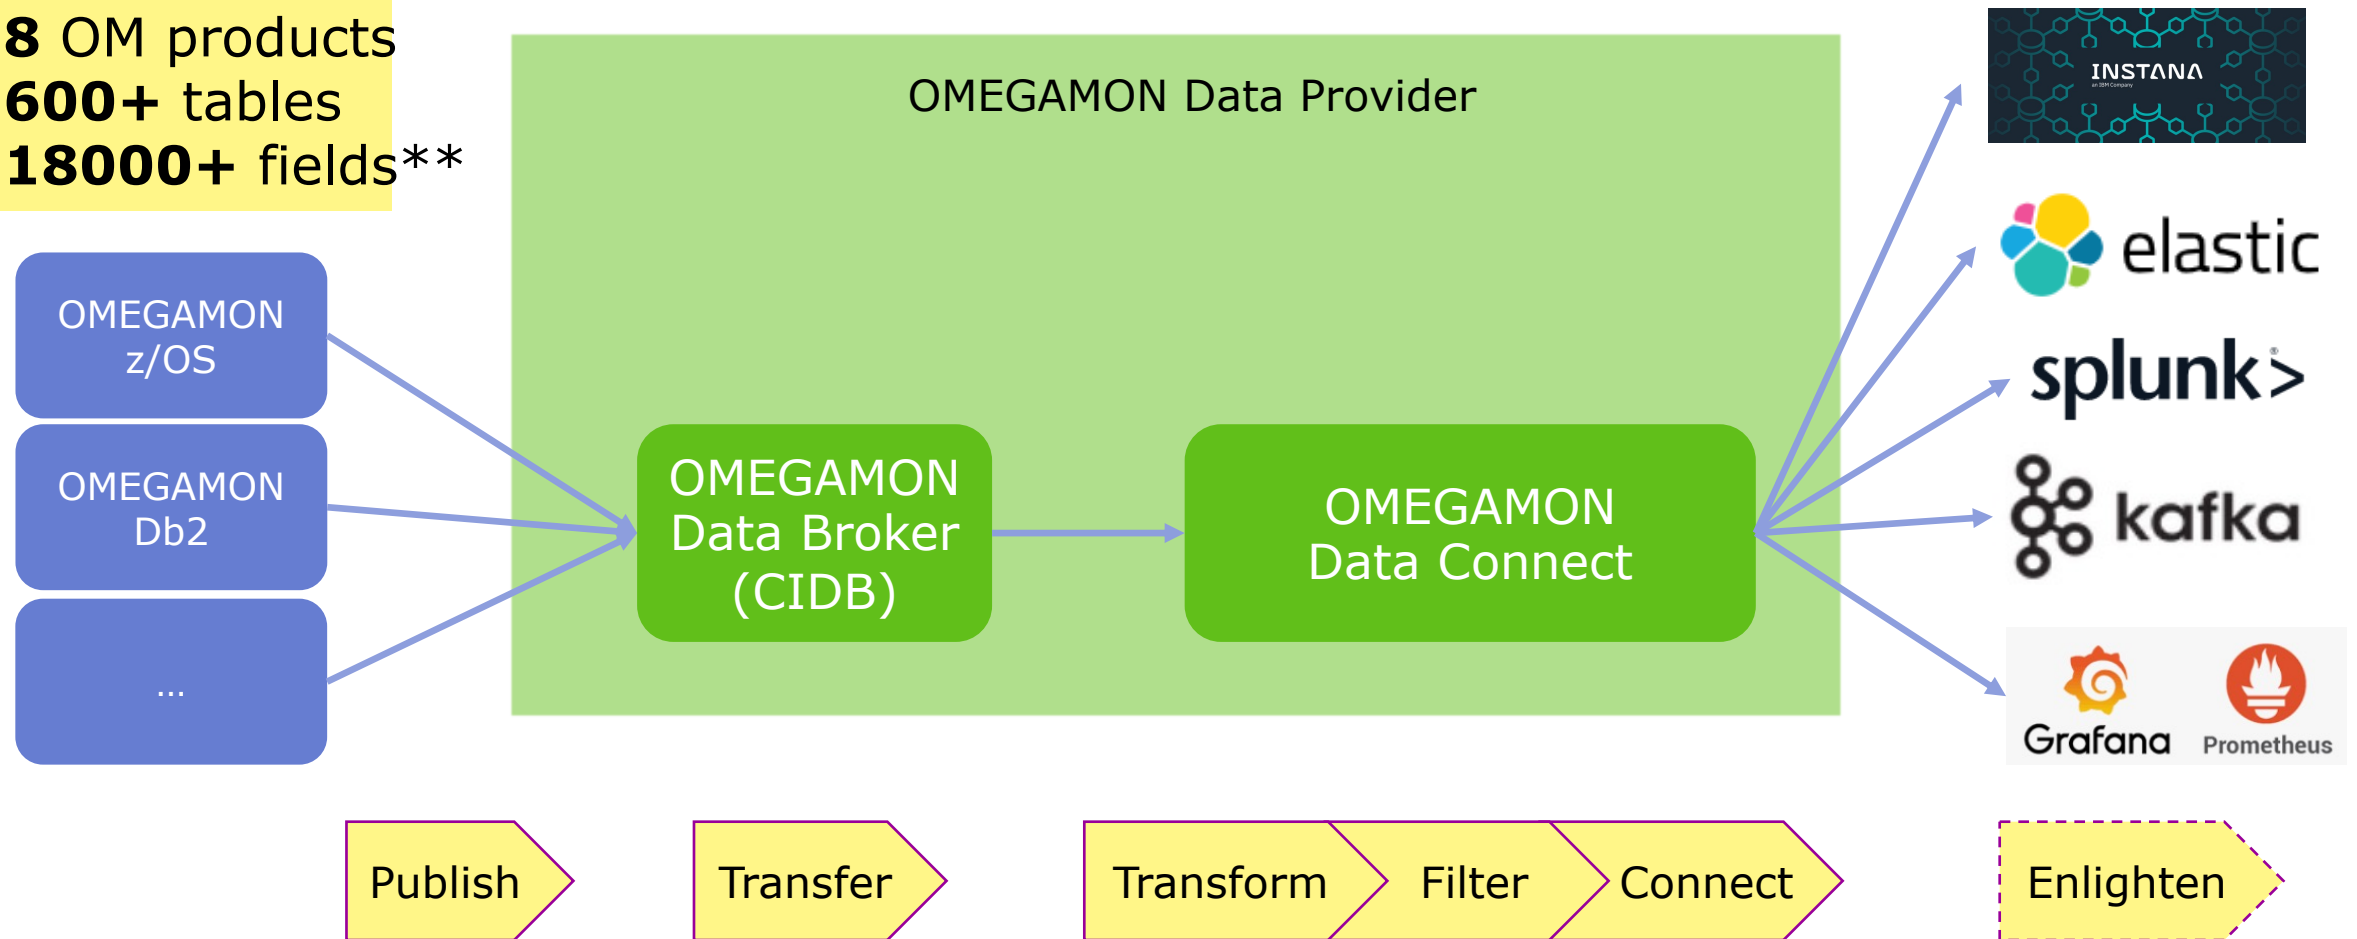
**OMEGAMON Data Provider (ODP)** serves OMEGAMON (attribute group) data to the destination or platform of your choice in a friendly and easily consumable format (JSON) – in near real-time

This provides a pathway for OMEGAMON to participate in the world of AIOPs –operational analytics and machine learning

Starter Elastic dashboards are provided to fast-track your take-up and implementation

# OMEGAMON Data Provider

**8 OM products**  
**600+ tables**  
**18000+ fields\*\***



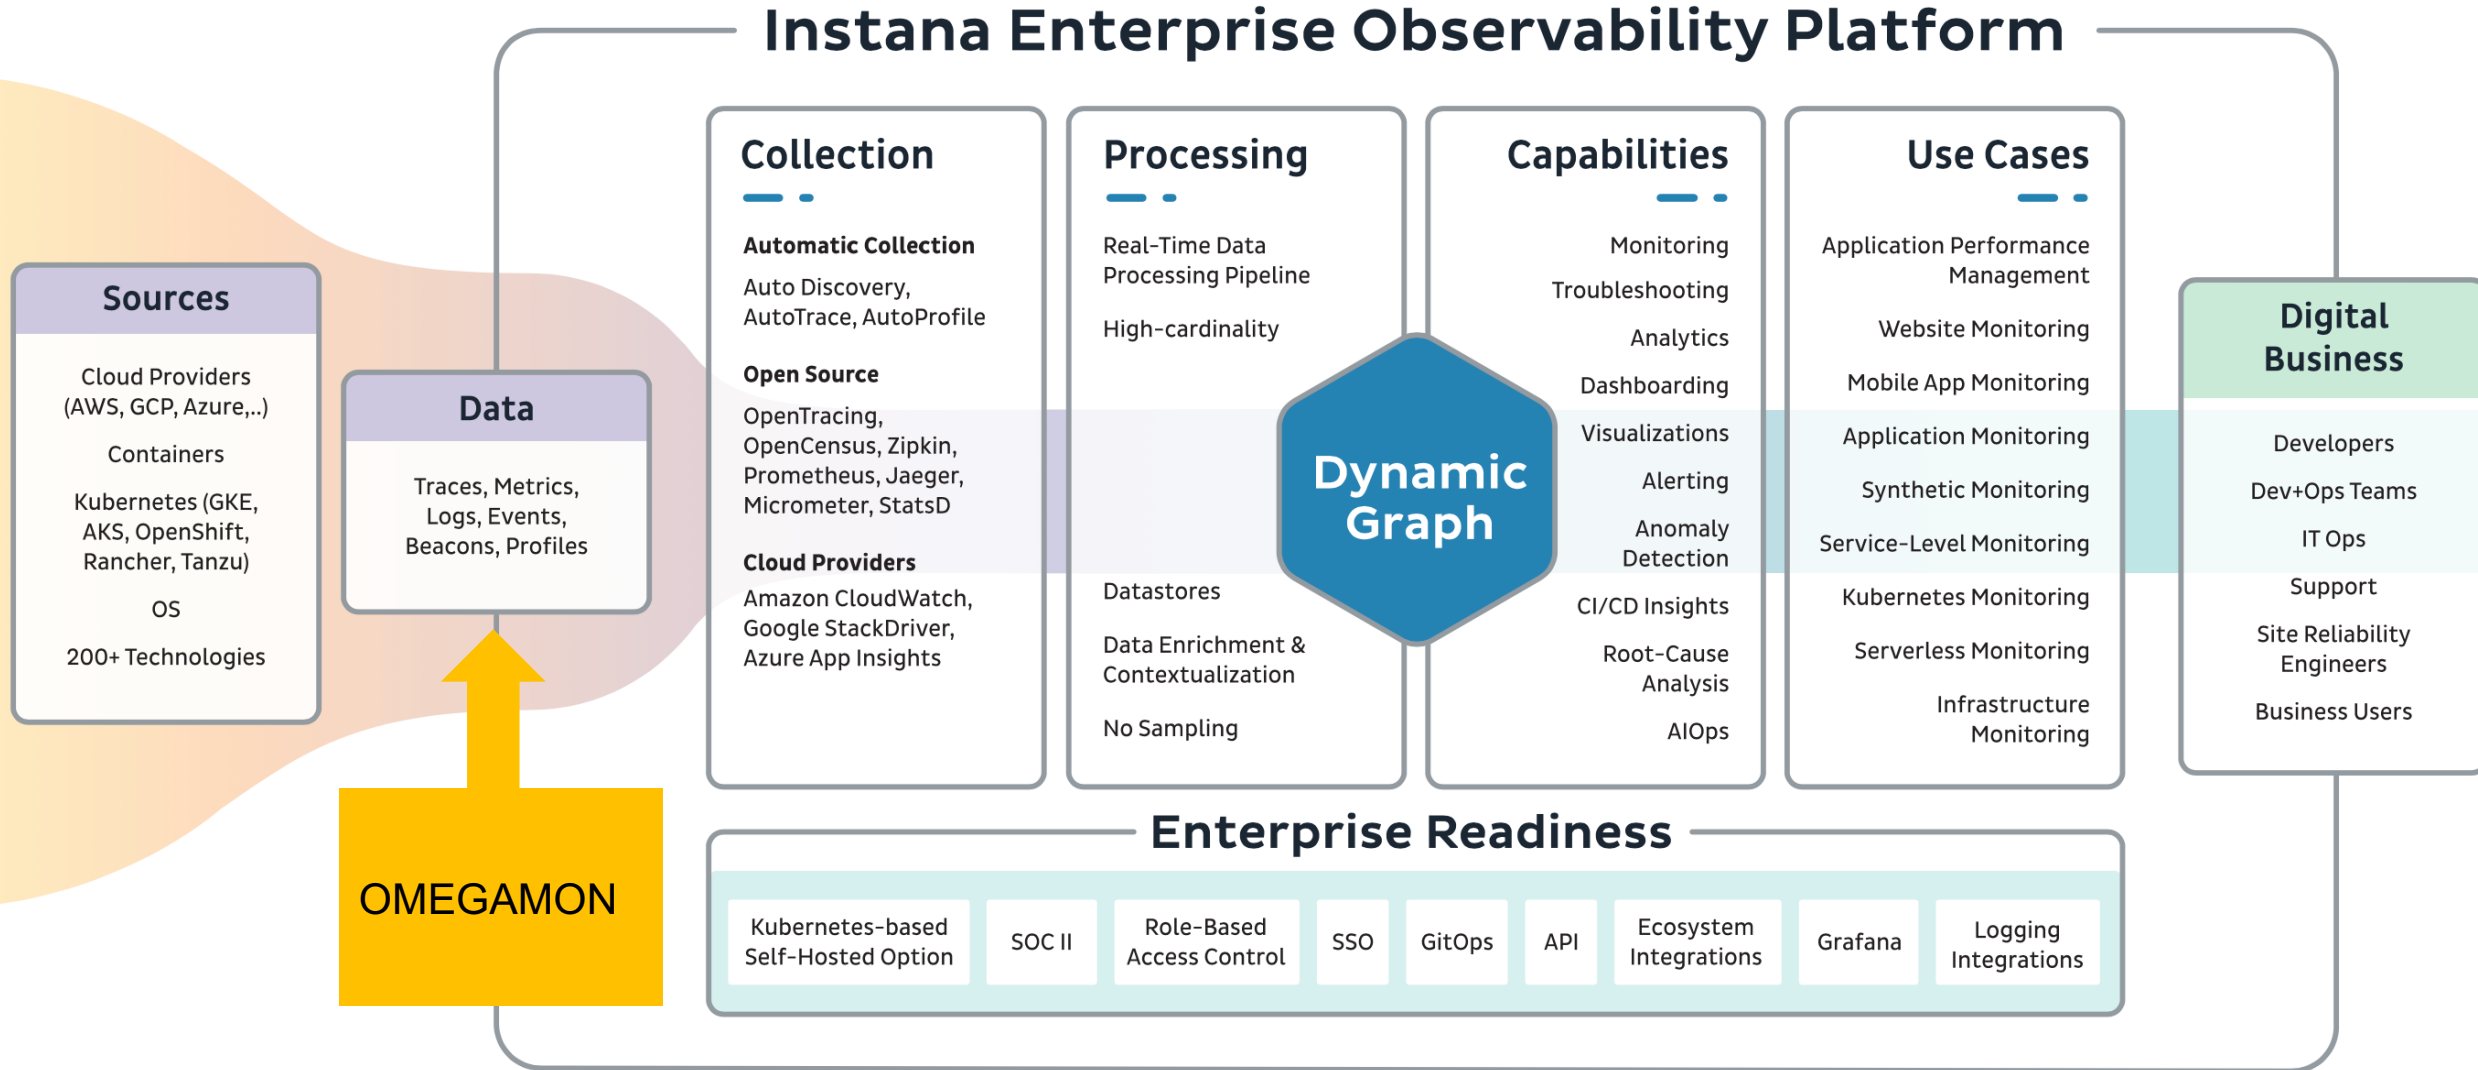
\*\* These are the number of metrics possible. If anyone actually streamed all of these to any of the analytics platforms listed on the right, we can guarantee your CPU would increase and therefore MSU's and software license charges as well. Best to curate a list of Key Performance Indicators to those targets, as has been done via a fixed list to Instana...roughly 30 tables with their associated metrics.

# Demo



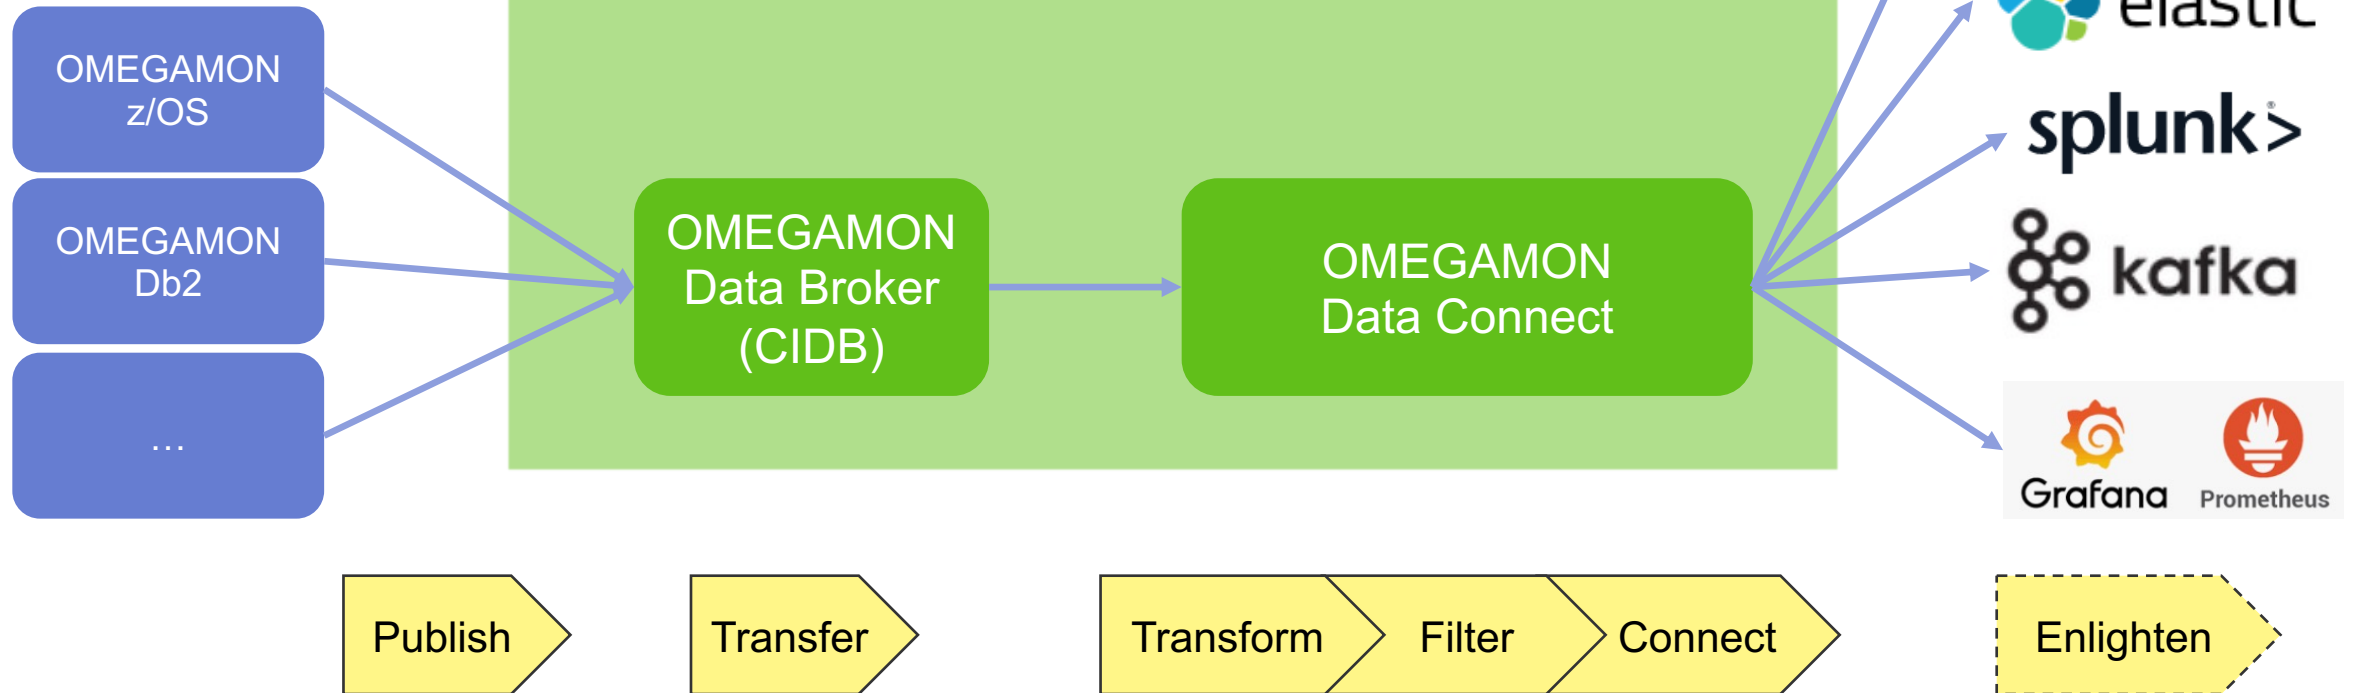
# Outsourcer and Large Enterprise Value

# Instana – overview

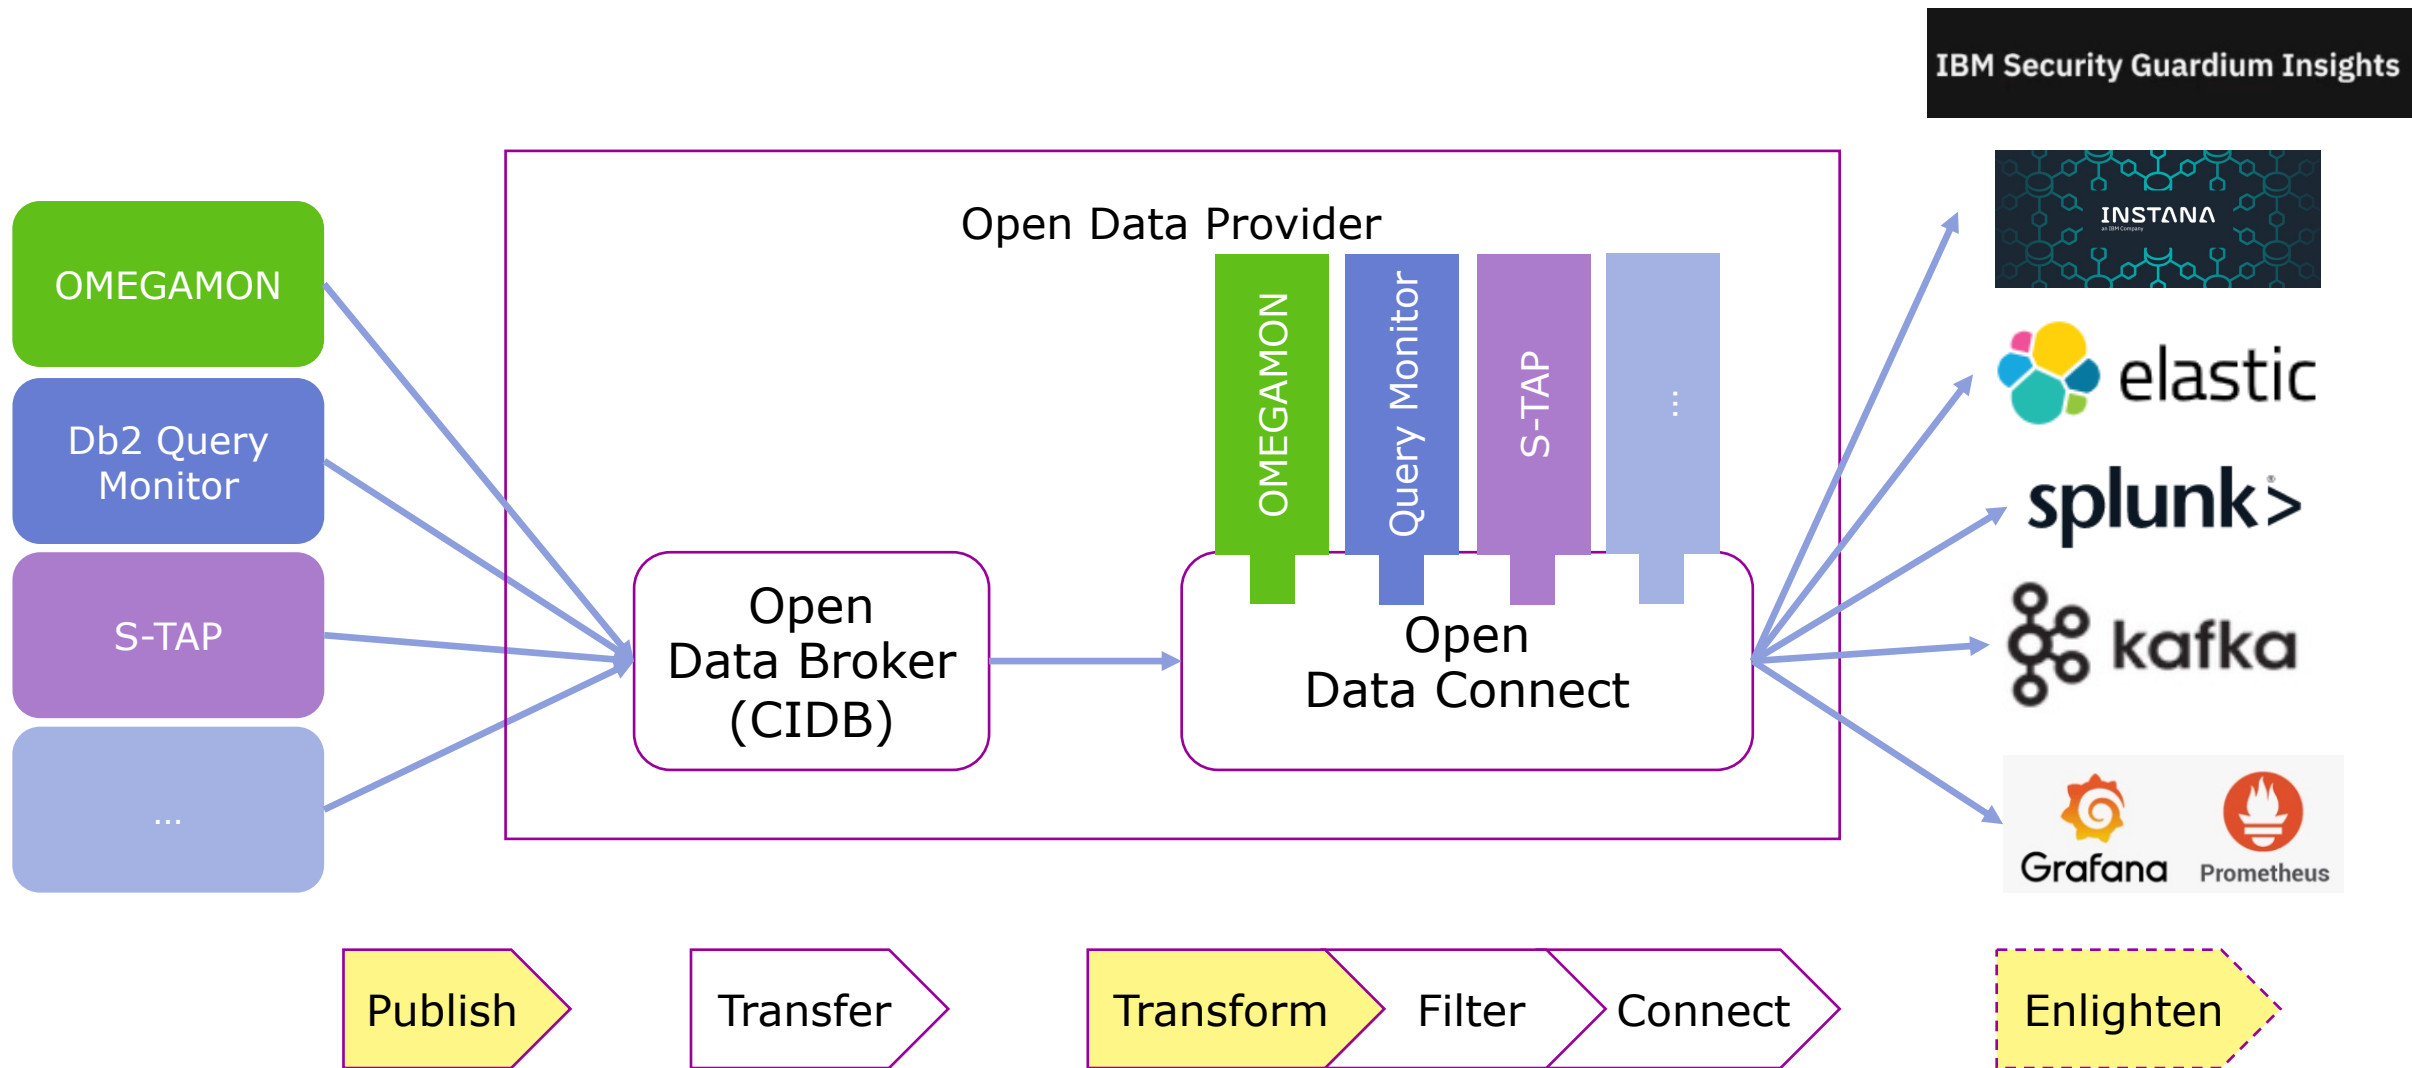


# OMEGAMON Data Provider

**8 OM products**  
**600+ tables**  
**18000+ fields**

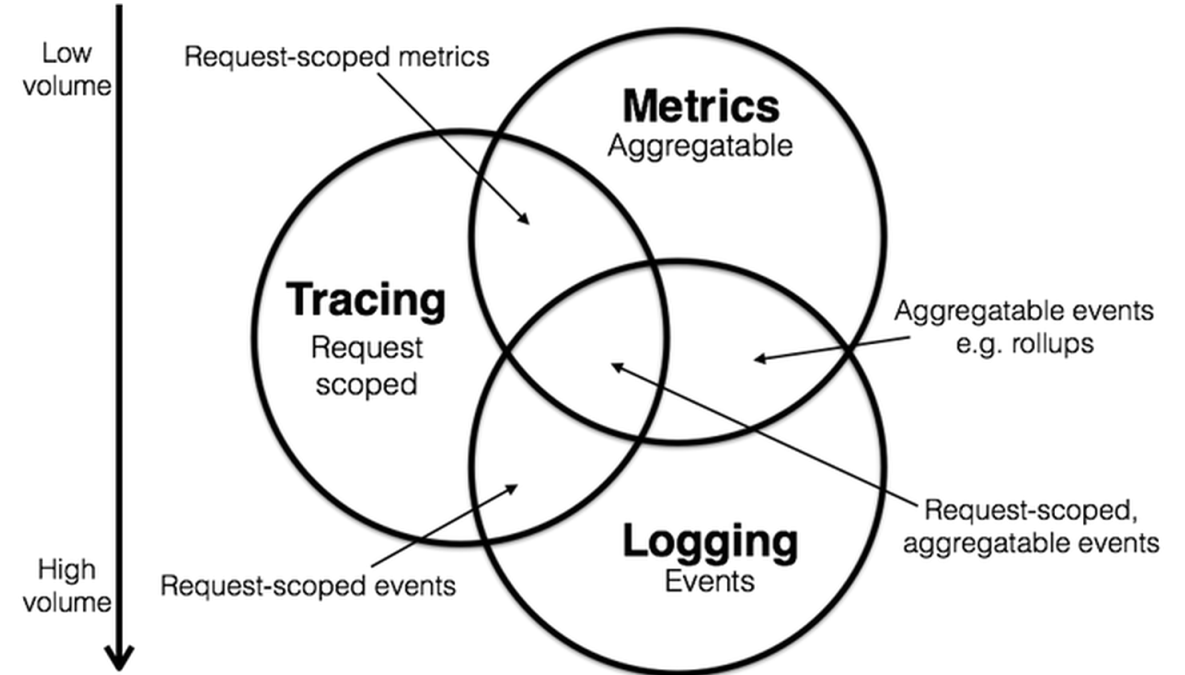


# Open Data Provider



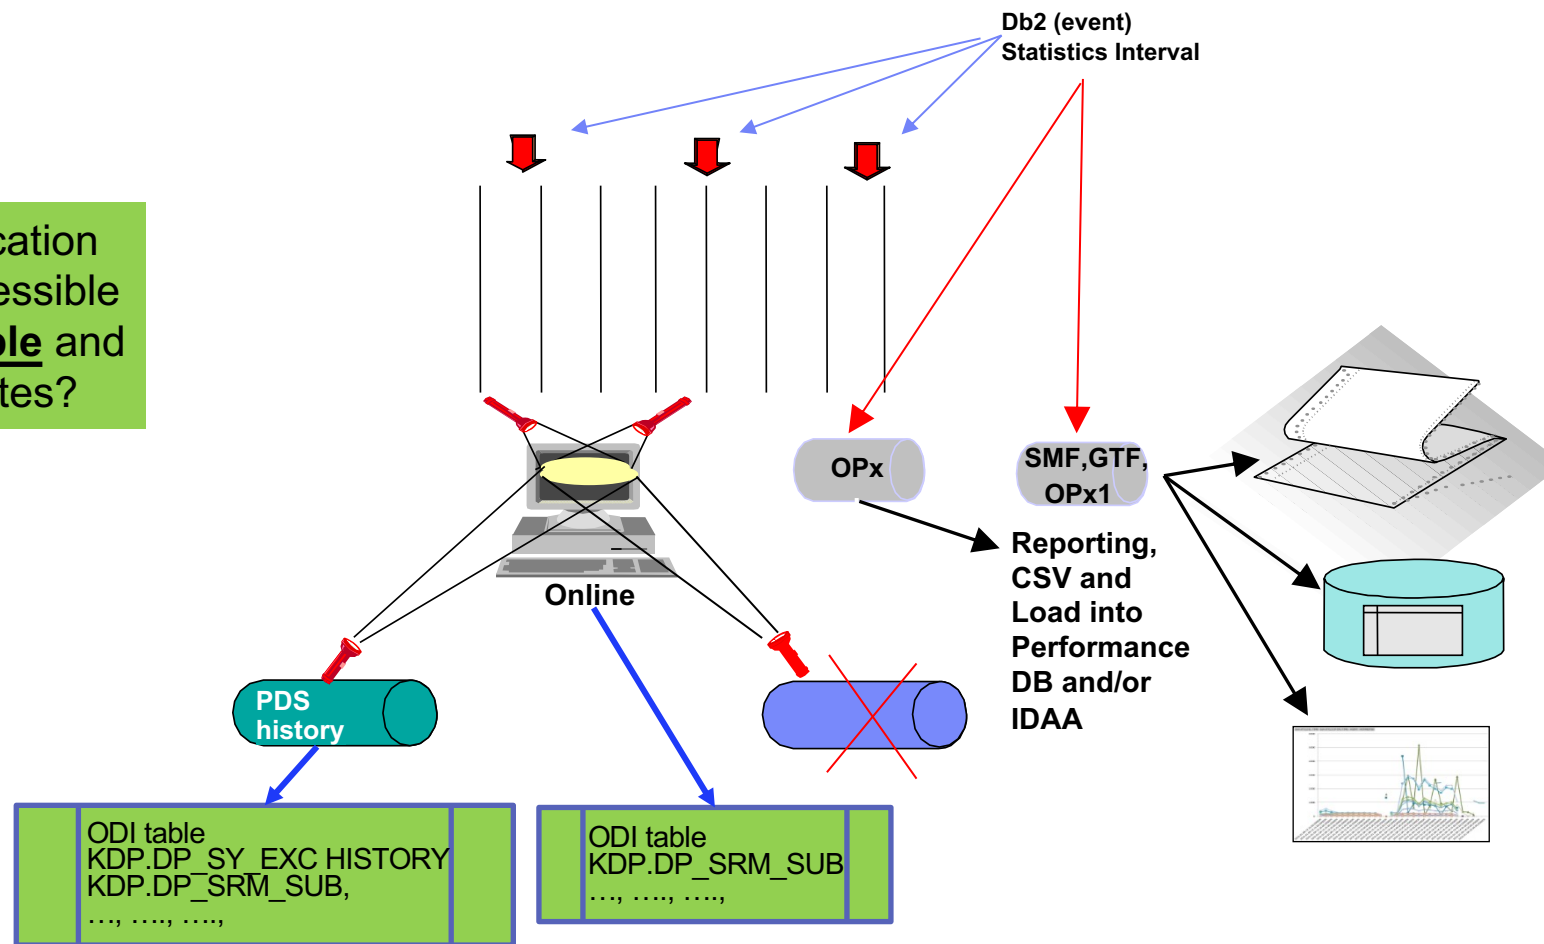
# Candidates for exploiting ODP SDK




- Db2 Query Monitor – work committed to IBM Data
- CICS, Db2, IMS Task History and Accounting (in progress)



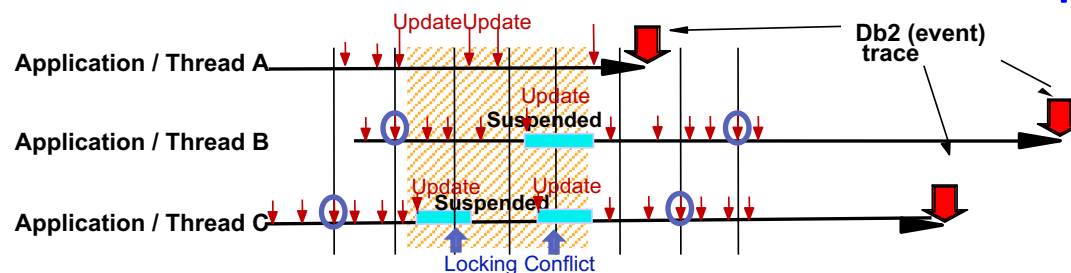
# Event trace data < versus > Real-time/snapshot data for Subsystem STATS

What application data is accessible via **ODI table** and their attributes?




- 
 Represents a Db2 Statistics trace record externalized at STATIME interval
-  Every 1 In the chart represents a snapshot of the status at this timestamp, e.g. every 5 second

# Event trace data < versus > Real-time/snapshot data for Applications

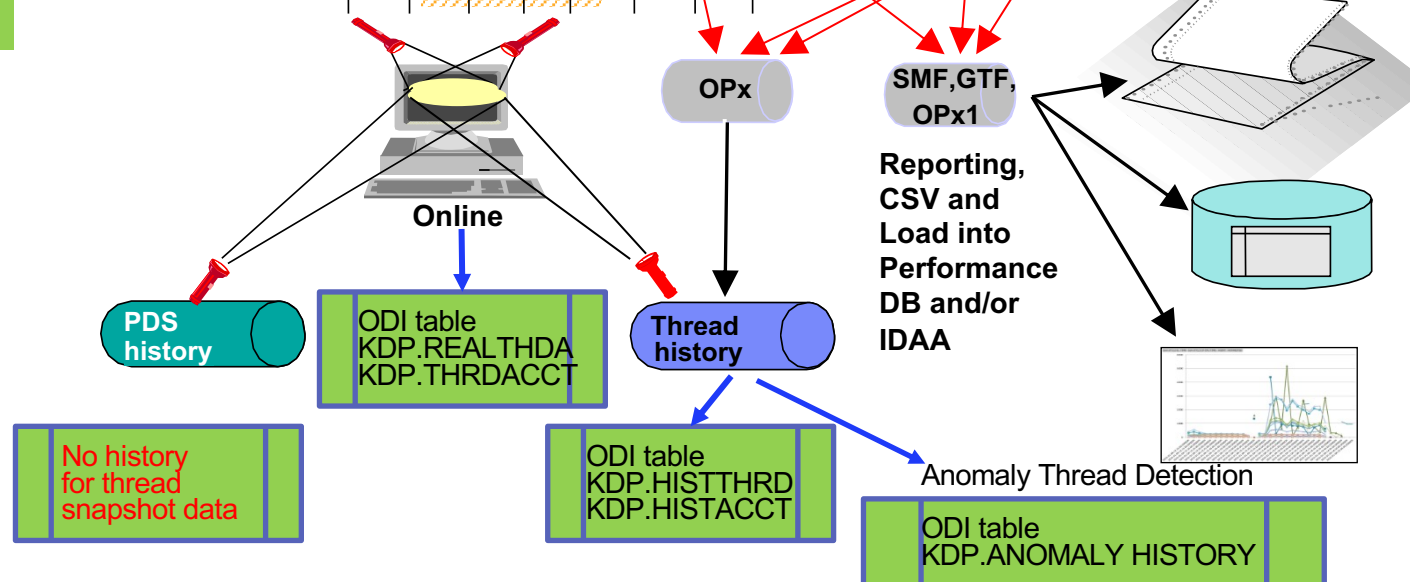


- ↓ → Represents a SQL call – e.g. between snapshot requests
- ⊙ → Represents a SQL captured at a snapshot request





- ↓ → Represents a completed / finished transaction
- Every I in the chart represents a snapshot of the status at this timestamp, e.g. every 5 second

What application data is accessible via **ODI table** and their attributes?








# Bridge the skills gap with modern user experiences ... optimized for the right user role




DevOps Team: Operator, Admin, Manager, SME, Developer, ...  
Chatbot

## IBM Z ChatOps

Collaborative problem resolution integrated in your chat platform



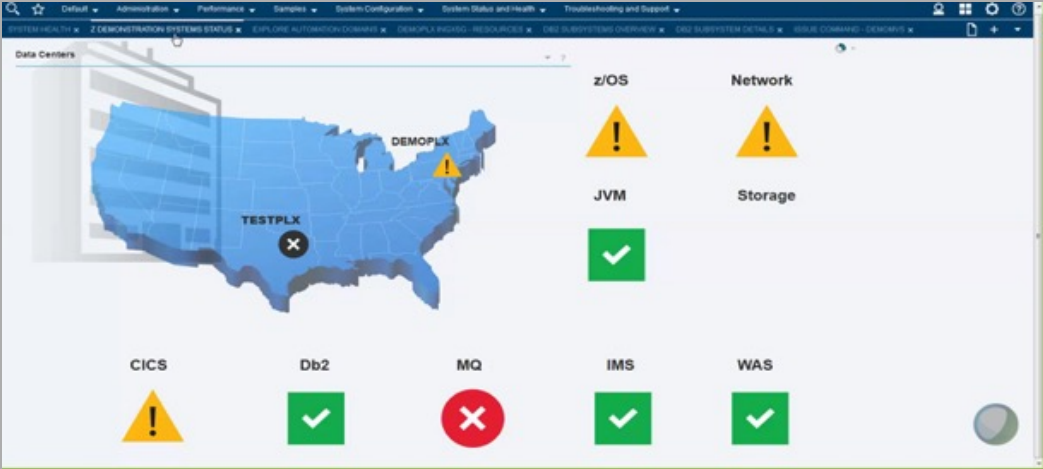
Slack MatterMost MS Teams

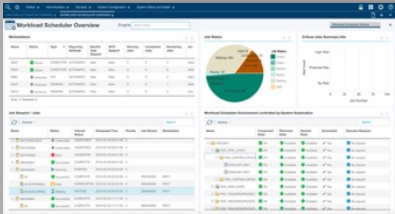


Z Operator


## IBM Service Management Unite

- Time-saving integrated web dashboards
- Brings all disciplines together
- Guided problem isolation & custom dashboards

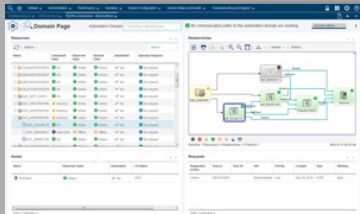





SMU Workload Scheduler:  
Dashboards for  
Z Workload Scheduler



SMU Performance Management:  
Dashboards for OMEGAMON




SMU Automation:  
Dashboards for Z System Automation and Z NetView




Z SME


## IBM Z Systems Management Tools




Anomaly Detection




Automation



Monitoring



Scheduling



Network Management

© Rocket Software, Inc. or its affiliates 1980 – 2022. All rights reserved. Rocket and the Rocket Software logo are registered trademarks of Rocket Software, Inc. Other product and service names might be trademarks of Rocket Software or its affiliates.

~~ Goal: Seamlessly plug into the broader experience ~~



# Where to find out more

Announcement Letter – November 9<sup>th</sup>, 2021: [Here](#)

## Product Documentation

- ODP Installation and User's Guide: [Here](#)
- Open z Data Demo content and guide on GitHub: [Here](#)
- Instana Observability for IBM z/OS: [Here](#)

## Blogs

- Introduction to ODP: [Here](#)
- Installation considerations for ODP: [Here](#)
- Now streaming CICS and Db2: [Here](#)
- CICS & Db2 Dashboards available: [Here](#)
- Usage Examples of ODP: [Here](#)
- Now streaming IMS and JVM: [Here](#)

## Video

- 8+ minute overview video: [Here](#)



**Master Blog that points to all of these (*bookmark this one*):** [Here](#) – This page is constantly updated



# Sponsor User Opportunities

Influence the tools you are using!

# Executive Summary

---

Machine Learning algorithms are being added to OMEGAMON work streams to address consumption, thresholds and response time

These will evolve over several functional updates, beginning with z/OS, Networks and JVM agents at the new V6.1 level:

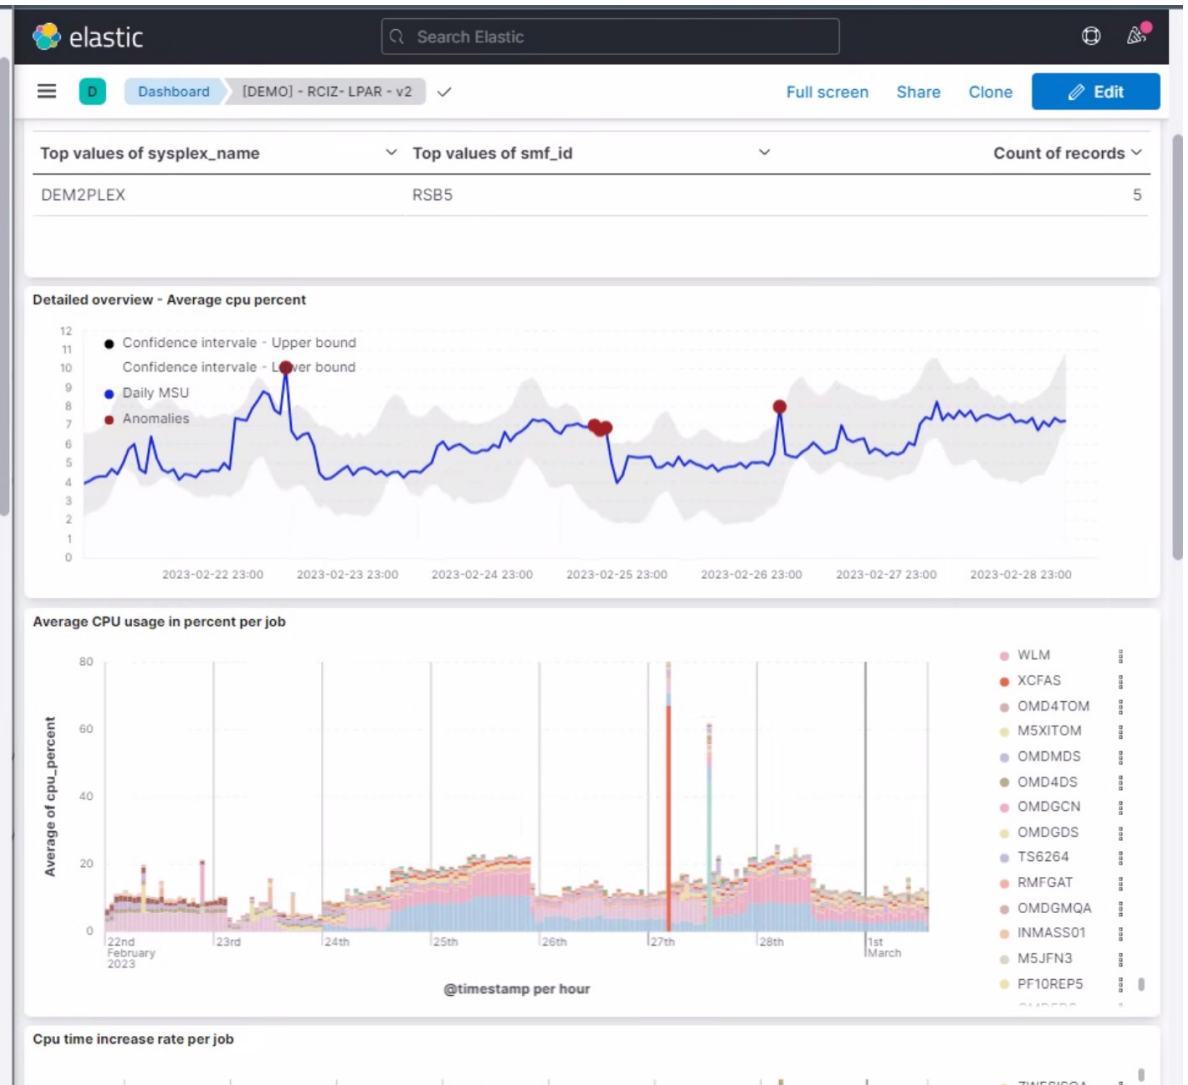
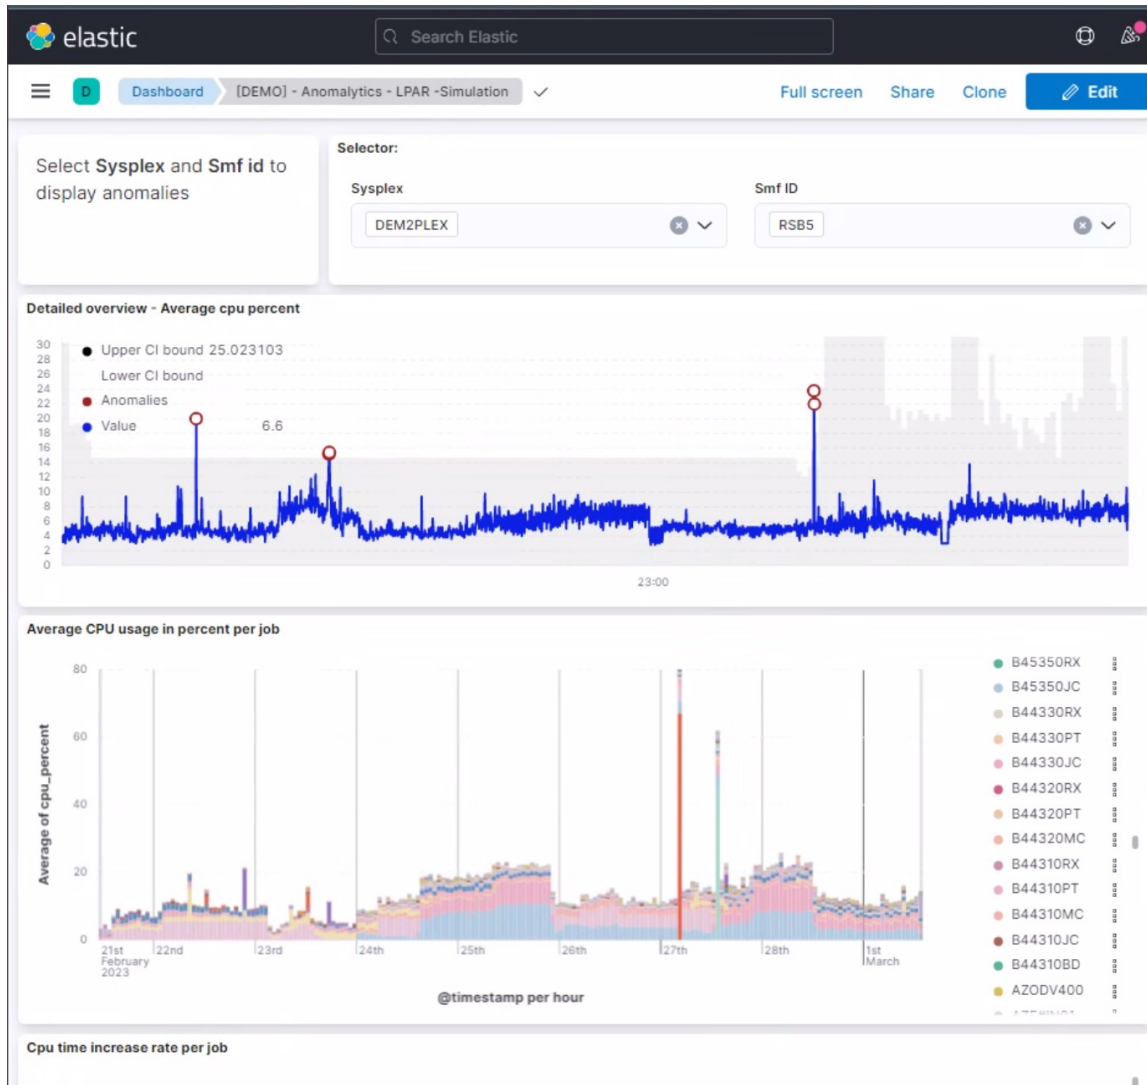
- KPI's for z/OS identify anomalies in MSU usage resulting in \$ Savings
- KPI's for Networks identify network anomalies that affect transaction volumes → \$'s lost
- KPI's for Java identify xxxx that affect yyyy that results in → zzz
- **Why is this different than situations? Complimentary**

Additional algorithms and agents will be updated in the future

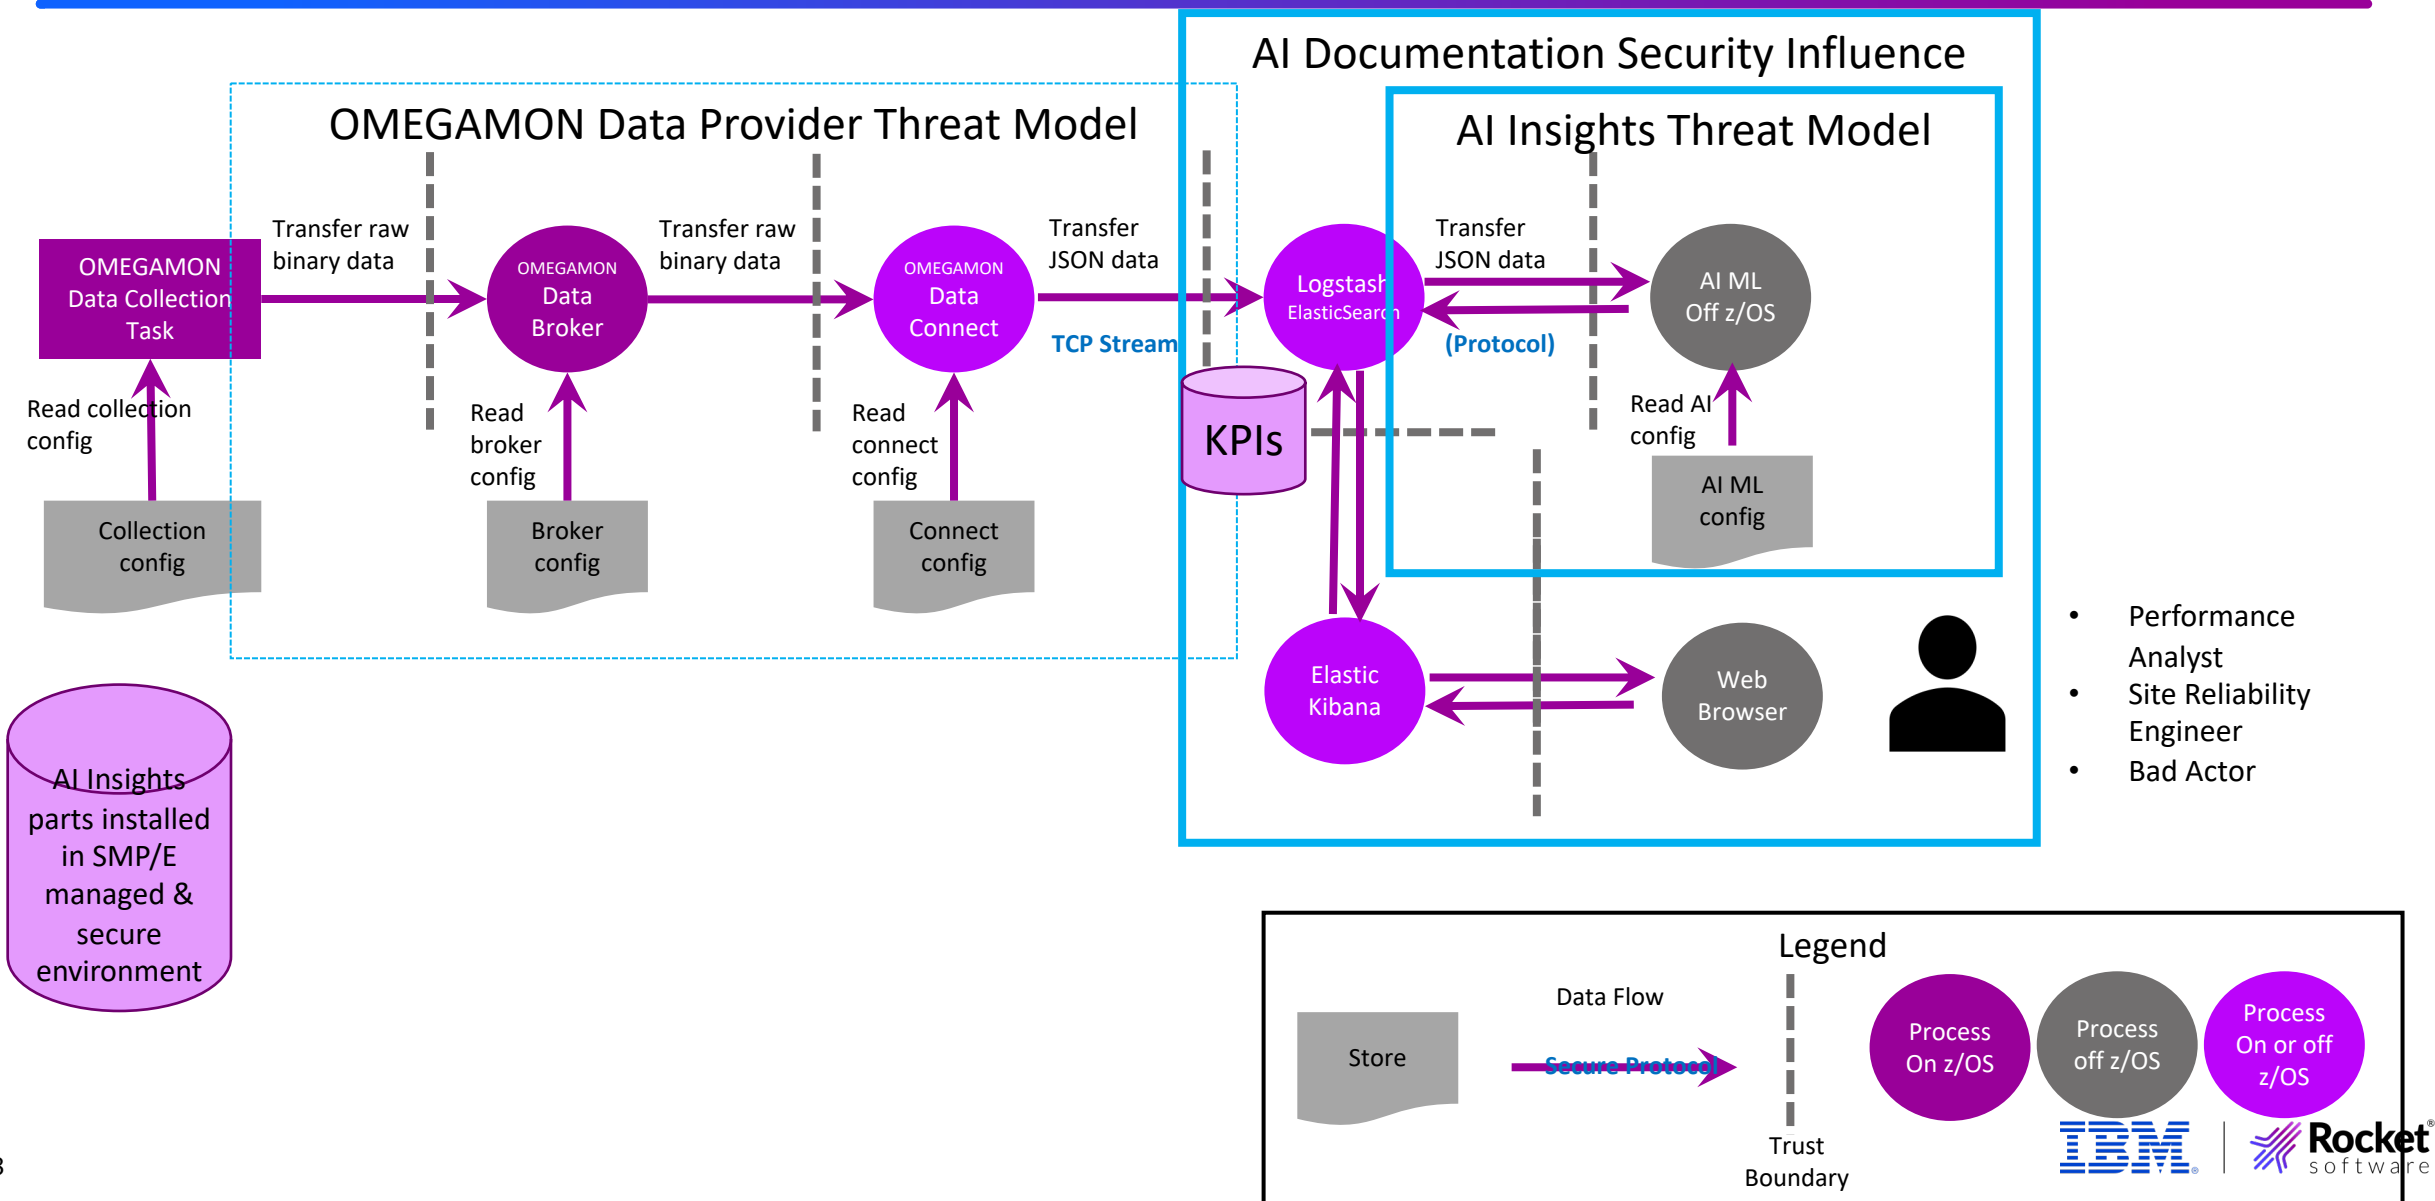
Customer data and customer situations would be good to get

# Under construction

# Two algorithms, similar anomalies



# OMEGAMON AI Insights Architecture – off z/OS



# Share with us the challenges you are facing ...

---

Do you face situations where ... ?

# Get ML driven insights into your CPU consumption

---

01

Identify a production SYSPLEX to start with

02

Upload 3 months of SMF 70,72 for the selected environment

03

SMF records injected in the ML Insight Factory

04

Presentation of the ML detected CPU consumption insights

# Your contacts for the SMF upload

---



Jim Porell

Principal Software  
Architect

---

[jporell@rocketsoftware.com](mailto:jporell@rocketsoftware.com)



Fabien Gautreault

Engineering  
Manager

---

[fgautreault@rocketsoftware.com](mailto:fgautreault@rocketsoftware.com)



Ludovic Granger

Product Owner

---

[lgranger@rocketsoftware.com](mailto:lgranger@rocketsoftware.com)

## Step by step SMF upload procedure:

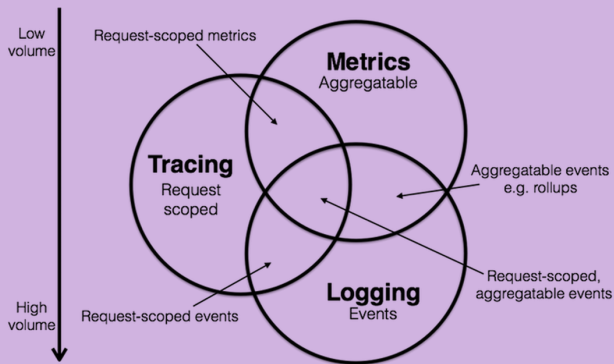
Offer valid until 3Q 2023



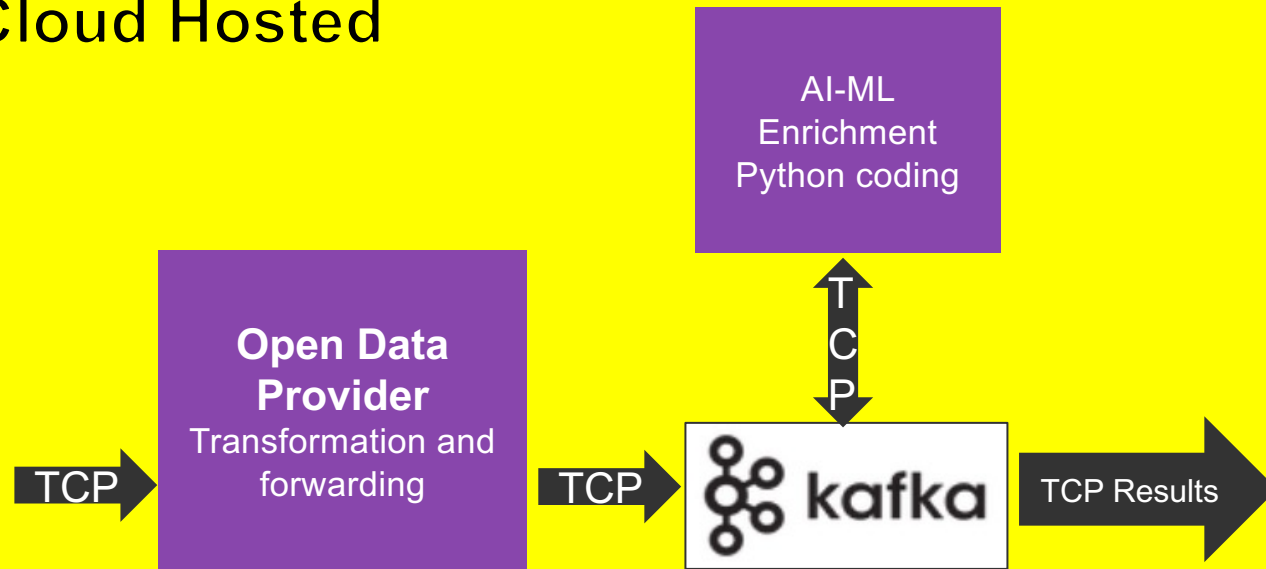


# AI-ML optimization Possibilities

## z/OS Hosted



## Cloud Hosted



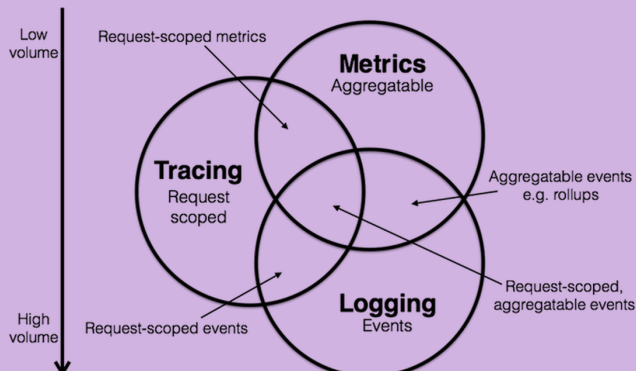
### Compare Models

- Elapsed Time?
- Deployment effort?
- TCO?

## Cloud Analytics

z/OS

## z16 Hosted



**Open Data Provider**  
Transformation and forwarding

AI-ML  
Enrichment  
Python coding  
**Telum Exploitation**

**kafka**

**Results**

## Cloud Analytics

z/OS

# Executive Summary

---

Customers have already proven to **benefit financially** and **reduce their time to resolve issues** through AI and ML

AI/ML is not a product...it's technology that can be applied to any product

OMEGAMON is embracing AI/ML to improve its handling of performance management


Customer input is needed to improve the training of models. **IBM and Rocket are looking for sponsor users for this new journey**

**I want you to enjoy this...nothing earth shattering except where the new technology can take us!**

# Questions?

# IBM can help in your AIOps journey

---



## Assessment

Qualify for a free  
AIOps Assessment


Contact your IBM  
Representative or  
Luke de Kansky  
[Lpdekans@us.ibm.com](mailto:Lpdekans@us.ibm.com)



## Handbook

Best practices for  
taking a hybrid  
approach to AIOps

<http://ibm.biz/AIOps-handbook>



## Collaborate

Register to gain  
access to Early  
Programs

[Client Feedback  
Program](#)



## Community

Join the IBM  
zSystems AIOps  
Community

<http://ibm.biz/AIOpsCommunity>

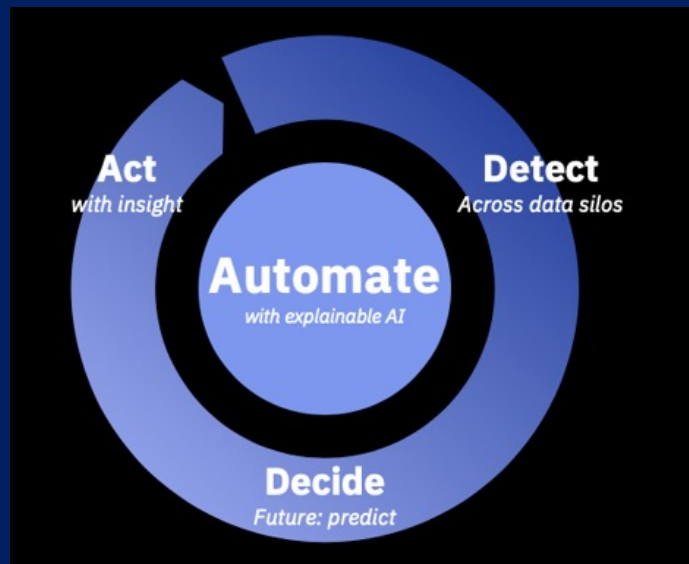


Join the Monitoring  
and Observability  
community

<http://ibm.biz/ZAIOPS-observability-blogs>

# AIOps for IBM zSystems

Improve systems management, IT operations, application performance and operational resiliency with Artificial Intelligence on the mainframe.

<https://ibm.biz/New-AIOps>



- **Detect:** monitor hybrid infrastructure and applications and detect incidents and anomalies.
- **Decide:** analyze issues and anomalies to isolate problems and identify root causes.
- **Act:** Rapidly remediate incidents to reduce impact on the customers and improved resiliency.

# Join Our Community of Experts!

<http://ibm.biz/AIOpsCommunity>

*Bringing together IT professionals to share their knowledge and expertise for leveraging AI-driven intelligence and IT Operations*




- Interact with subject matter experts and peers
- Community topic discussions
- Latest information - Blog postings
- Notifications of events
- Videos and Library content



<http://ibm.biz/AIOps-handbook>



<https://ibm.biz/AIOps-Video1>



<https://ibm.biz/Video-Auto-Res>

# Bonus Material Hints and Tips



# Where to find out more

Announcement Letter – November 9<sup>th</sup>, 2021: [Here](#)

## Product Documentation

- ODP Installation and User's Guide: [Here](#)
- Open z Data Demo content and guide on GitHub: [Here](#)
- Instana Observability for IBM z/OS: [Here](#)

## Blogs

- Introduction to ODP: [Here](#)
- Installation considerations for ODP: [Here](#)
- Now streaming CICS and Db2: [Here](#)
- CICS & Db2 Dashboards available: [Here](#)
- Usage Examples of ODP: [Here](#)
- Now streaming IMS and JVM: [Here](#)

## Video

- 8+ minute overview video: [Here](#)



**Master Blog that points to all of these (*bookmark this one*):** [Here](#) – This page is constantly updated

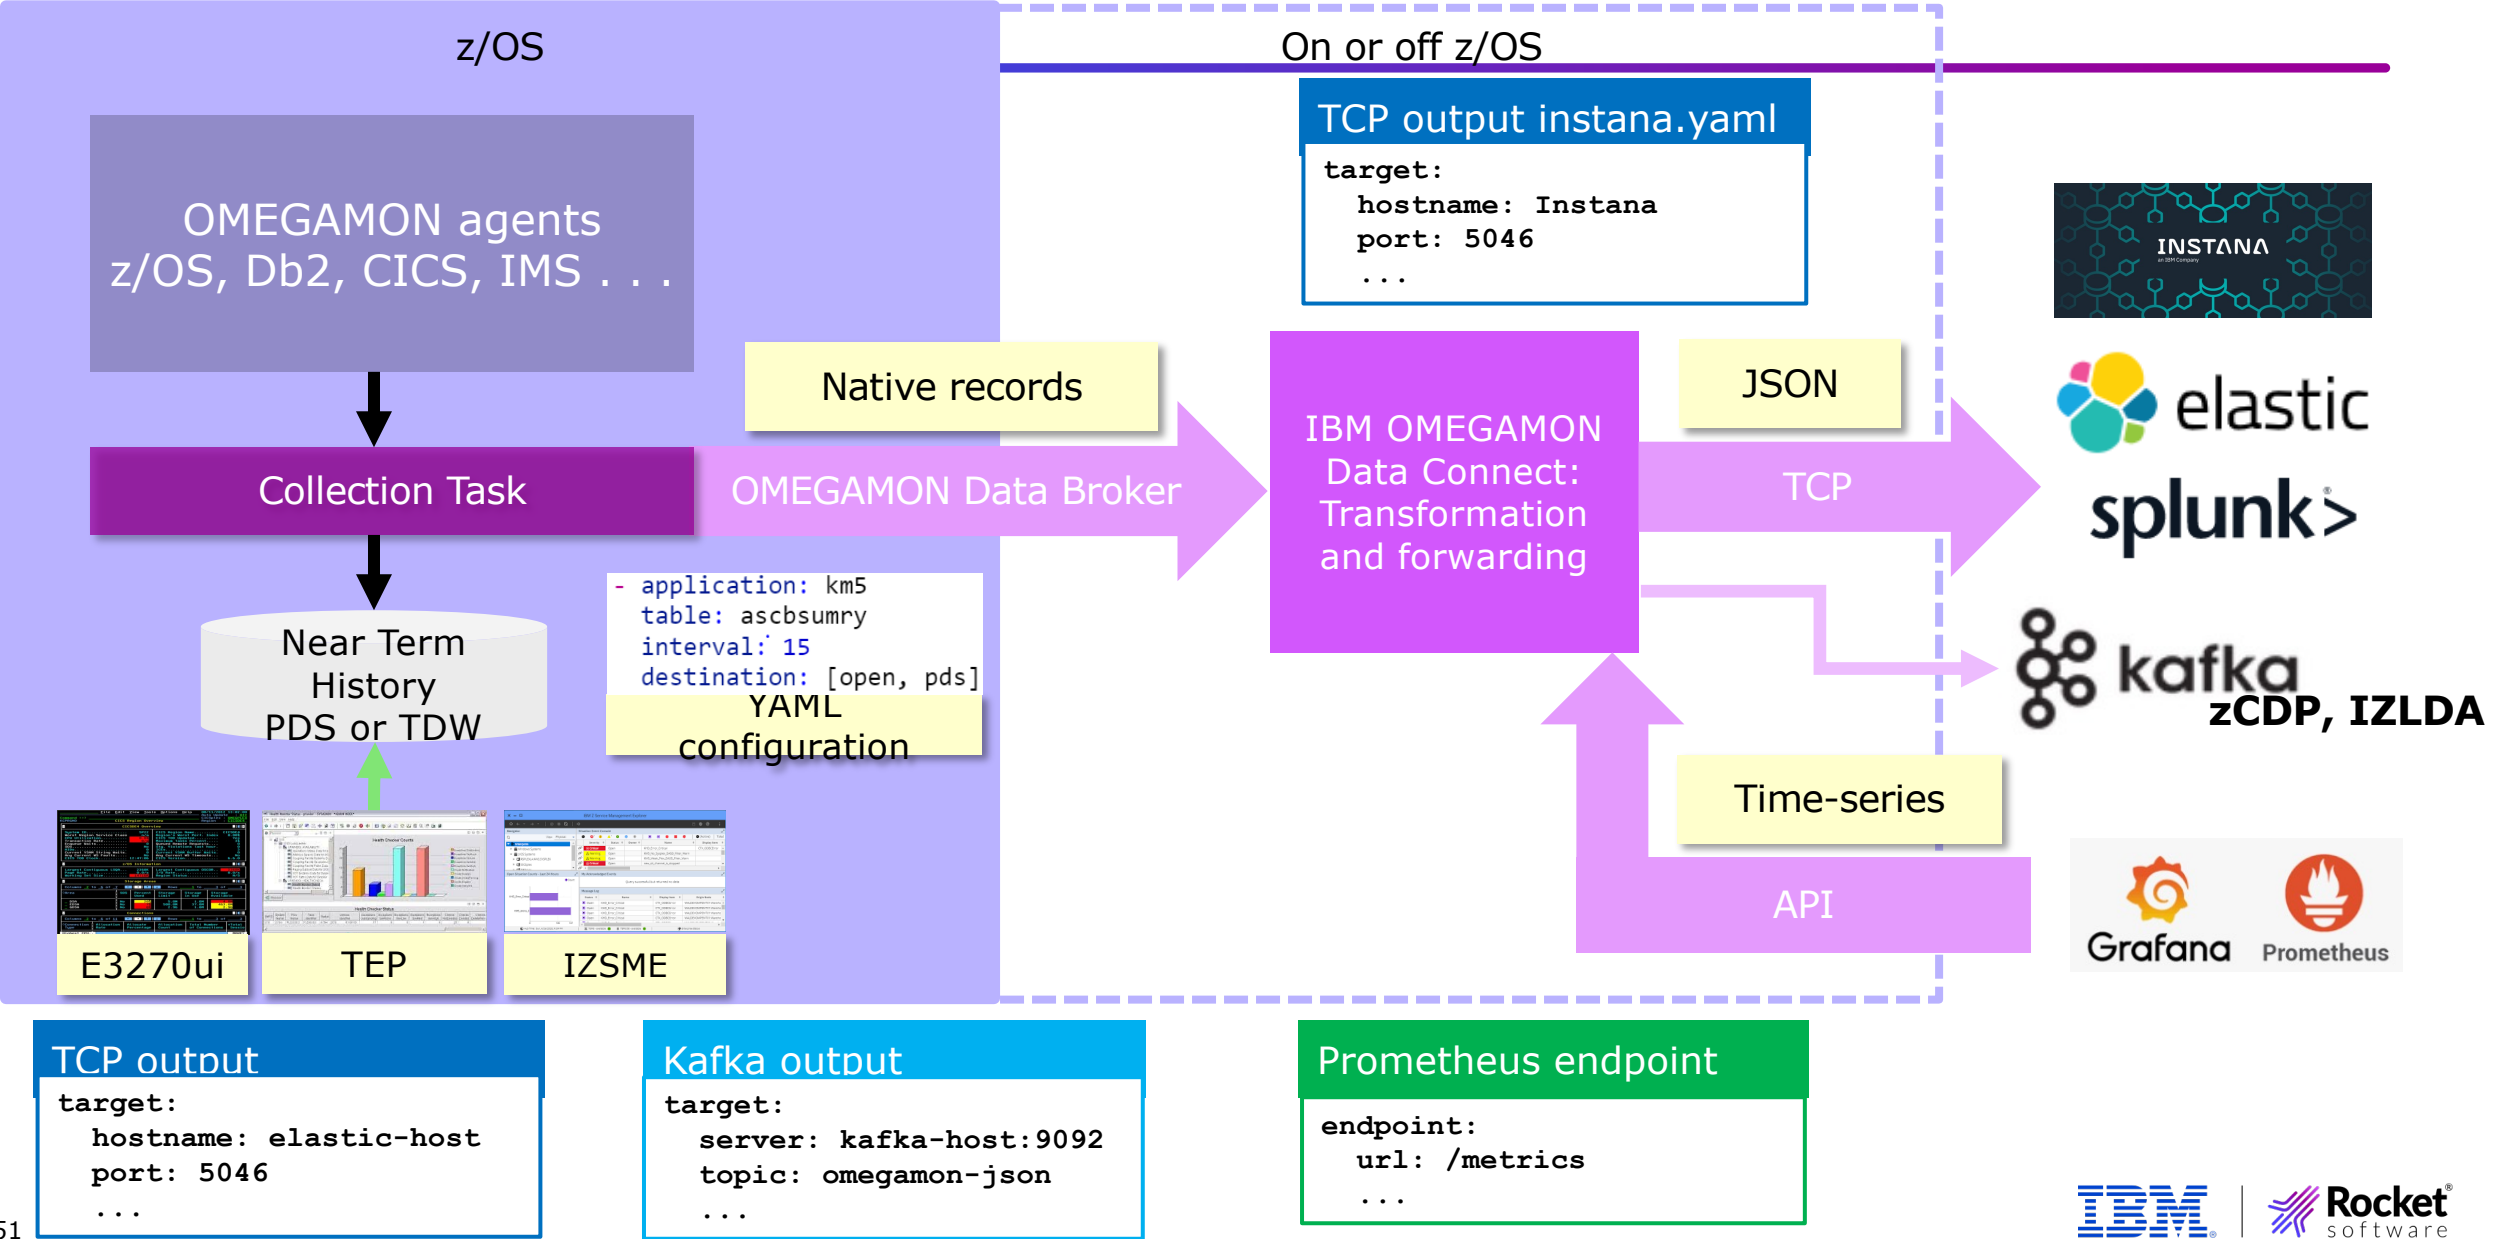




# Overview

- Background needed: Some working knowledge of the components of OMEGAMON Data Provider (ODP)
- Key areas requiring attention during the installation and customization
- A variety of hints and tips to avoid errors
- Most importantly, anything in this presentation constitutes a very high-level view. Please read the ODP Installation and User's Guide to get the details

# Architecture – overview



# Get the prereq code installed



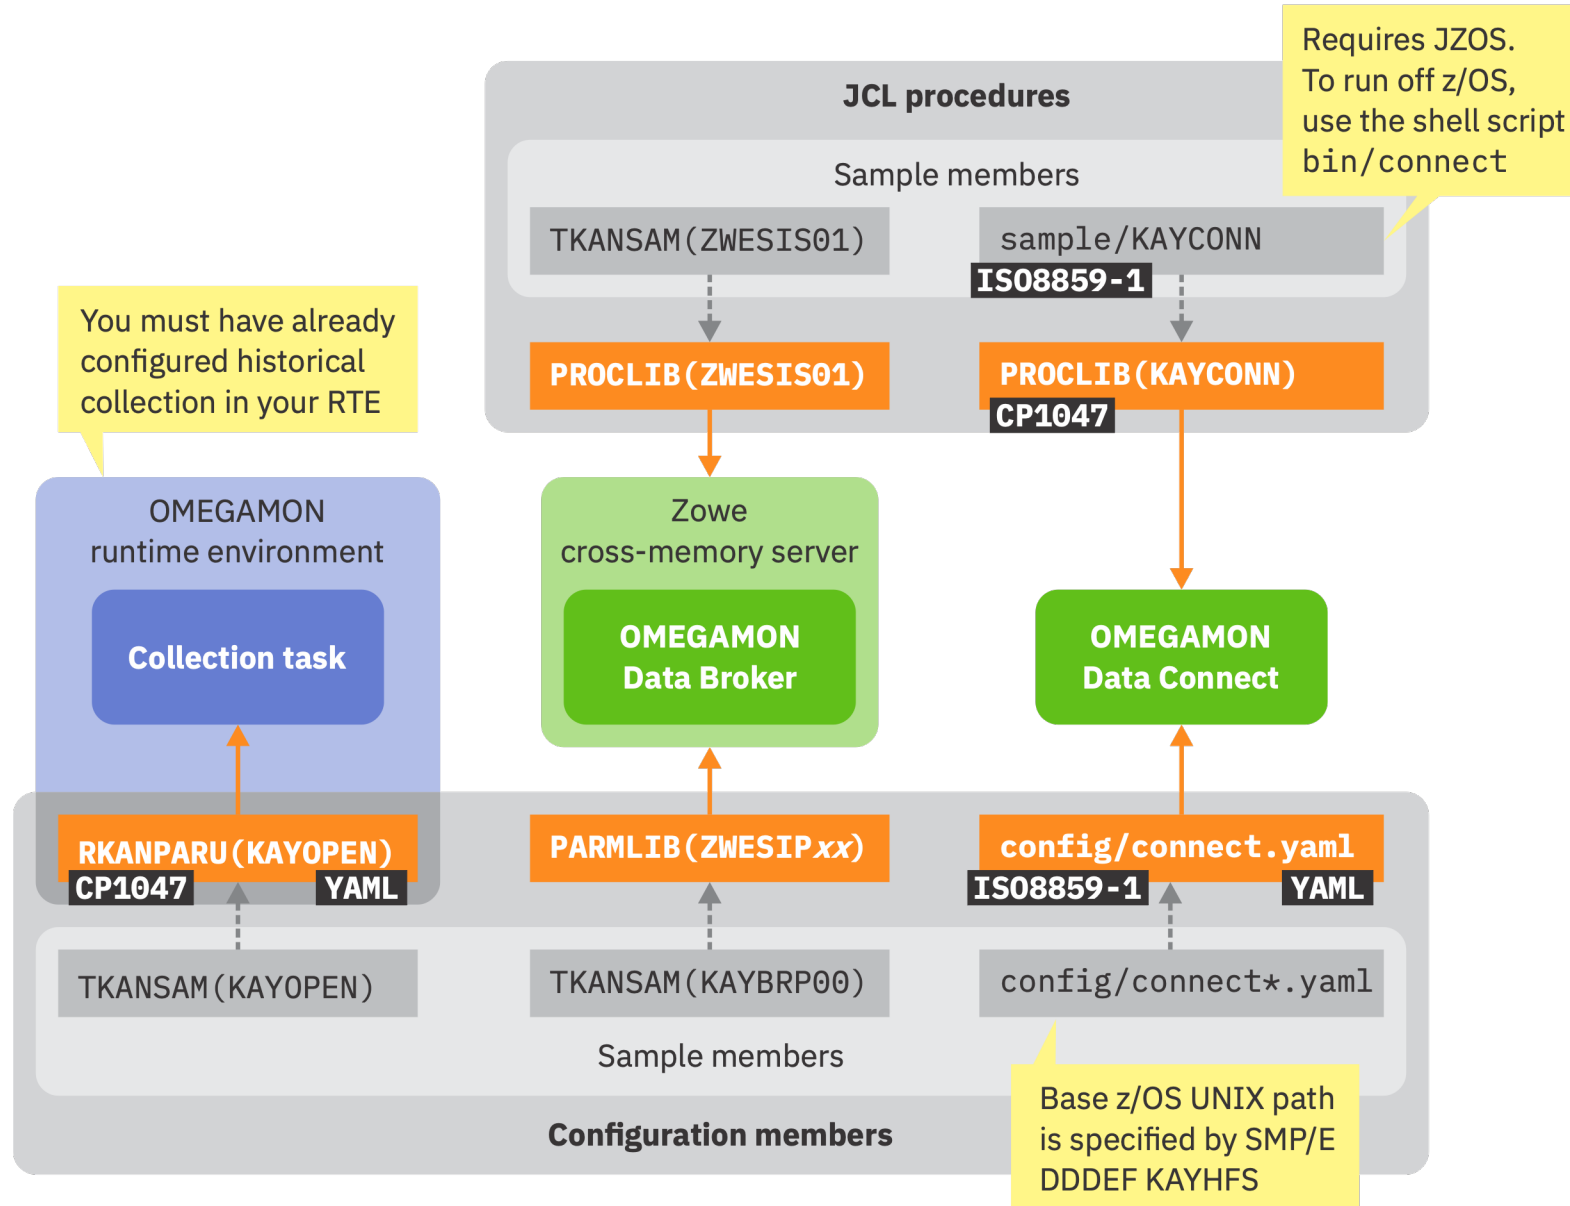
You need one of the following product suites:

# Decide on the architecture that is going to be deployed

---

Where will the ODP Data  
Connect Server run - on  
z/OS or off z/OS?

# Overview of Configurable parts



ODP has two types of configurable parts:

1. Members that **run** components:

- **PROCLIB(ZWESIS01)** JCL procedure that runs the Zowe cross-memory server that hosts OMEGAMON Data Broker.
- **PROCLIB(KAYCONN)** JCL procedure that runs OMEGAMON Data Connect. Alternatively, if you decide to run OMEGAMON Data Connect off z/OS, you can use the shell script bin/connect.

2. Members that **configure** components:

- **RKANPARU(KAYOPEN)** A YAML document that specifies collection configuration parameters, such as which attribute groups to send to OMEGAMON Data Broker.
- **PARMLIB(ZWESIPxx)** A plain-text member that specifies OMEGAMON Data Broker parameters, such as the host and port on which OMEGAMON Data Connect is listening.
- **config/connect.yaml** A YAML document that specifies OMEGAMON Data Connect configuration parameters, such as the output method for attributes

Network Security also can be defined.

- That will be described elsewhere



# Tip: Editing files – Correct Code page is important

---

Files within the z/OS UNIX file system can be tagged with a specific code page.

It's important that as you edit these files, the correct code page is utilized.

Default/rule of thumb for services running within z/OS native functions is CP1047 (EBCDIC)

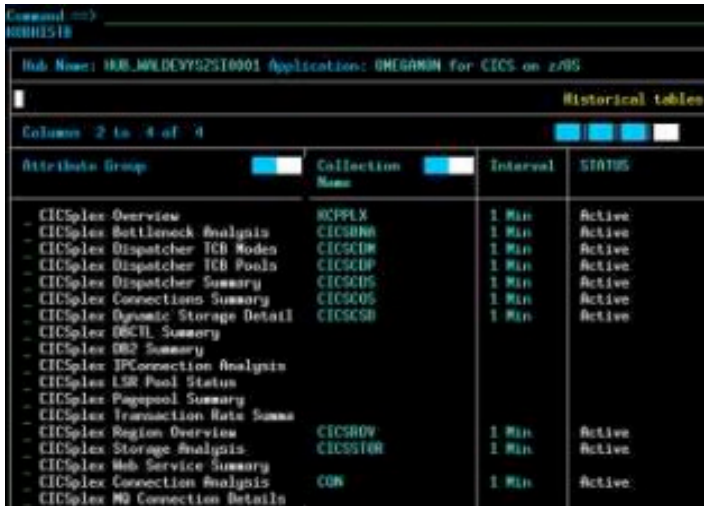
Default/rule of thumb for servers running within Java, Unix System Services and non-z is ISO8859-1 (ASCII)

Why is this important? Within YAML, the bracket symbols '[' ']' use different character tagging depending on the code page. We've seen cases where edits were done with CP 285 (UK style) that caused some difficulties. Please set up your 3270 terminal emulator or file editor properly so that these characters get translated properly within JCL and YAML.

# History Collection activation lifecycle

## OMEGAMON agent

- The OMEGAMON agents specify their attribute tables to be eligible for History Collection
- Done in E3270ui (TOM) or TEP
- Attribute tables are now eligible for ODP forwarding to analytics
- Not all attribute tables are eligible for collection. This is an agent decision.



The screenshot shows a terminal window with the command 'OMEGAMON' entered. Below the command, there is a table titled 'Historical tables' with columns: Attribute Group, Collection Name, Interval, and STATUS. The table lists various attribute groups and their corresponding collection names, intervals, and statuses.

Attribute Group	Collection Name	Interval	STATUS
CICSplex Overview	KCPPLX	1 Min	Active
CICSplex Bottleneck Analysis	CICSBBN	1 Min	Active
CICSplex Dispatcher TCB Modes	CICSDDM	1 Min	Active
CICSplex Dispatcher TCB Pools	CICSDDP	1 Min	Active
CICSplex Dispatcher Summary	CICSDDS	1 Min	Active
CICSplex Connections Summary	CICSDDS	1 Min	Active
CICSplex Dynamic Storage Detail	CICSDDC	1 Min	Active
CICSplex DBCTL Summary			
CICSplex DB2 Summary			
CICSplex IPCorrelation Analysis			
CICSplex LSR Pool Status			
CICSplex Pagepool Summary			
CICSplex Transaction Rate Summary			
CICSplex Region Overview	CICSROV	1 Min	Active
CICSplex Storage Analysis	CICSSTOR	1 Min	Active
CICSplex Web Service Summary			
CICSplex Connection Analysis	CON	1 Min	Active
CICSplex MQ Connection Details			

## Collection → Data Broker

KAYOPEN.yml contains the list of tables that are eligible to be sent to analytics servers

- Identify name of the Omegamon Data Broker. ZWESIS\_STD is default
- Identify the data to be streamed by:
  - Agent/product id
  - Attribute group table name
  - Recording interval (minutes)
  - Where to send (PDS or ODP/open)

broker:

name: ZWESIS\_IRF

collections:

- product: kc5
- table: kcpplx
- interval: 0
- destination: [pds, open]
- product: kc5
- table: tran
- . . .

## Data Connect

Connect.yml describes:


- Input – where the data comes from
- Output – where data is sent to
- Format – List of fields in the payload

```
input:
  tcp:
    host-name: acme_prod
    port: 15351
output:
  tcp:
    logstash:
      host-name: elastic.acmecorp.com
      port: 15046
filter:
  products:
    kc5:
      tables:
        kcpplx:
          fields:
            - cicsplex_name
            - number_of_regions
            - transaction_rate
            - io_rate
            - page_rate
            - storage_violations
            - current_string_waits
            - current_buffer_waits
```



# Identifying the Data to Collect from OMEGAMON

## Where does it fit



```
broker:
  name: ZWESIS_STD 1
collections:
  - product: km5 2
    table: assumry
    interval: 1
    destination: [pds,open]
  - product: km5
    table: ascpuutil
    interval: 1
    destination:
      - open
      - pds
```

## Update YAML file

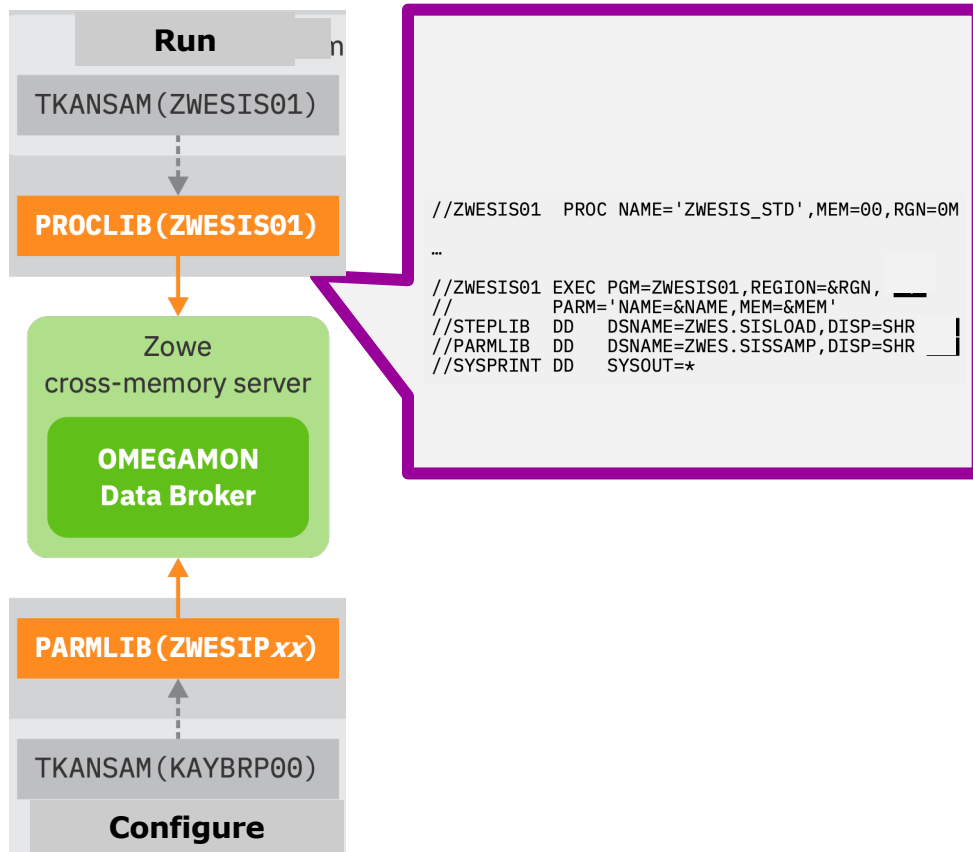
1. Identify name of Zowe cross-memory server. ZWESIS\_STD is default
2. Identify the data to be streamed by:
  - Agent/product id
  - Attribute table name
  - How often/interval (seconds)
  - Where to save/destination (PDS or ODP/open)

## TIPS

- Turn on history collection for the agents you desire in the TEP or TOM
- If you use the included Zowe cross-memory server, make sure it is copied and named correctly.
- Don't put this in the SMP/E version. Why? When you refresh or upgrade maintenance, your changes will be lost
- Use the correct code page

# Capturing the data from OMEGAMON

## Where does it fit



## Configure (and copy) ODP Data Broker plugin

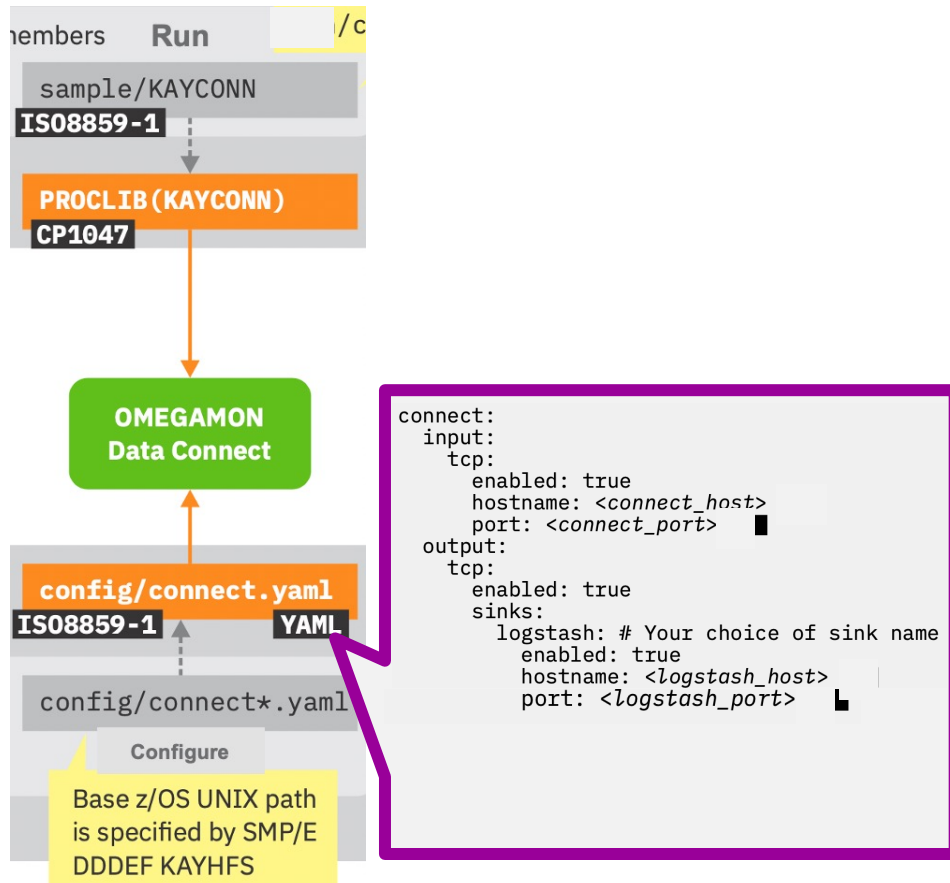
1. Update PARMLIB for ODP plugin to Zowe cross-memory server
  - Update IP @ and port of ODP Connect server
2. Set up Zowe cross-memory server if not in place already
  - PPT
  - APF
  - JCL to PROCLIB
  - Started Task

## TIPS

- STEPLIB containing ODP broker must be PDSE
- Don't (re)start Zowe cross-memory server until all components configured
- Don't put this in the SMP/E version. Why? When you refresh or upgrade maintenance, your changes will be lost
- Use the correct code page
- Follow the steps in the manual

# Making the data available to Analytics servers

## Where does it fit



## Configure (and copy) ODP Data Connect server

1. Z/OS: Configure Started task. Non-z: Set up script
2. Update connect.yaml
  - Input: ip@ of this server instance
  - Output: ip@ of analytics target
  - Identify any field filtering for collected tables
3. See the documentation to configure analytics servers

## TIPS

- Don't (re)start ODP Connect server until all components configured
- Don't put this in the SMP/E version. Why? When you refresh or upgrade maintenance, your changes will be lost
- Use the correct code page
- Follow the steps in the manual

# Tip: Configuring the Analytics servers

Elastic: Requires Logstash config and Elasticsearch Index template

# Other Configuration considerations

---

# Lessons Learned

---

Prometheus setup around metrics is numbers, not strings

- Labels are text/strings

While Data Connect can handle multiple sinks/personas in a single instance, for production, best to dedicate a Data Connect server to single sink

Code page syntax issues with [ ] square bracket characters

- Brackets not part of invariant EBCDIC character set. Terminal emulator, ISPF and file need to coordinate code pages. Difficult with EMEA national codepages set on 3270 emulator (e.g TE: 285, ISPF: 1047, connect.yaml: iso 8859-1)
- If necessary, change [open,pds] to:
  - open
  - pds

YAML syntax (indentation) is critical. Use free tool: <https://jsonformatter.org/yaml-validator>

- Have customers share yaml files, not cut and paste into email

Attribute Names are case sensitive. CICS\_Region\_name is wrong. cics\_region\_name is correct

If data doesn't appear to be streaming (i.e. no KPQH038I with attribute table name not found in log), check KAYOPEN yaml and TOM/TEP historical collection set up


- After changing either yaml file, refresh Data Broker or Data Connect to pick up change

# Outsourced or Very Large Customers

## Customer Personas

- Operator
- Subject Matter Expert
- Admin
- Executive

```
1 broker:
2   name: OM-prod-broker
3
4 collection:
5   - product: km5
6     table: jobsumry
7     interval: 1
8     destination: [open, pds]
9     fields:
10    - lpar: {EQ: [SYS1, SYS2, SYS3]}
11    - cputime: {GT: 1.0}
12    - jobname
13    - program
14
15 target:
16   - name: splunk
17     hostname: 10.117.198.102
18     port: 15046
```



## ODP Configuration Connect.yaml

Include file for each Persona

- Greatly simplifies size of yaml file
- Can direct Personas to different sinks
  - Executive to Splunk; others to Elastics

Dev/test:

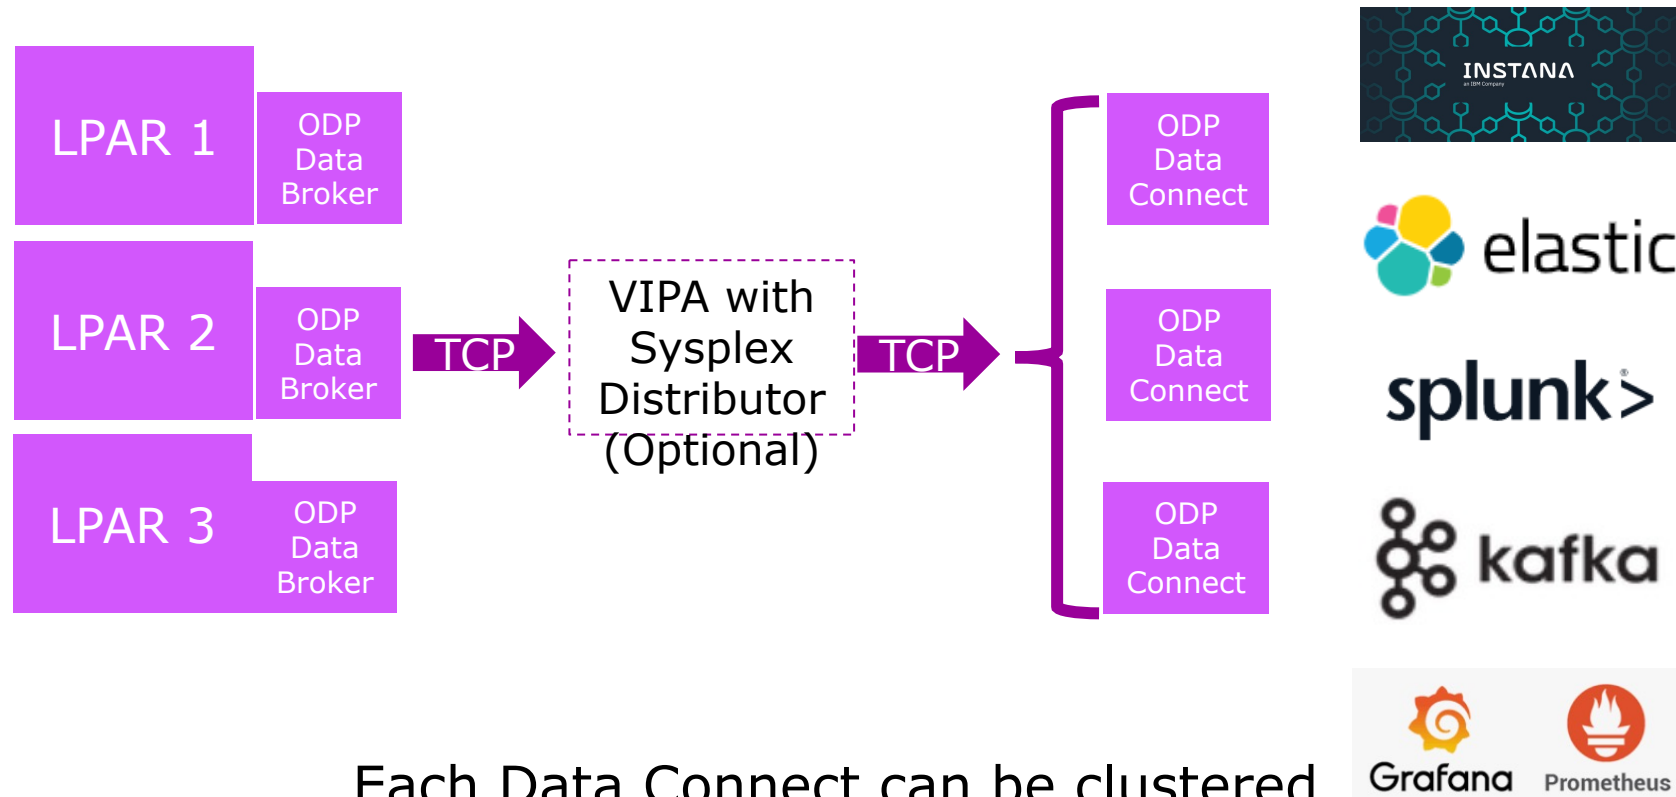
- Single Data Connect with multiple includes
- Production
- One or more Data Connect with single include
- Dedicated to one or more personas

Filtering

- Field level filtering to further distinguish personas
  - Only product LPARs for Customer A
  - Only CICS jobs for SME A:CICS
- Field level filtering for alerts
  - Only stream data above a metrics threshold



# Clustering for segregation and resilience



Each Data Connect can be clustered  
Each "sink" can be clustered  
Customers can be isolated to their own clusters

# Scalability & Redundancy

**This topic did not make it into the Installation and User's Guide, but is an important consideration as customers move toward a production environment**

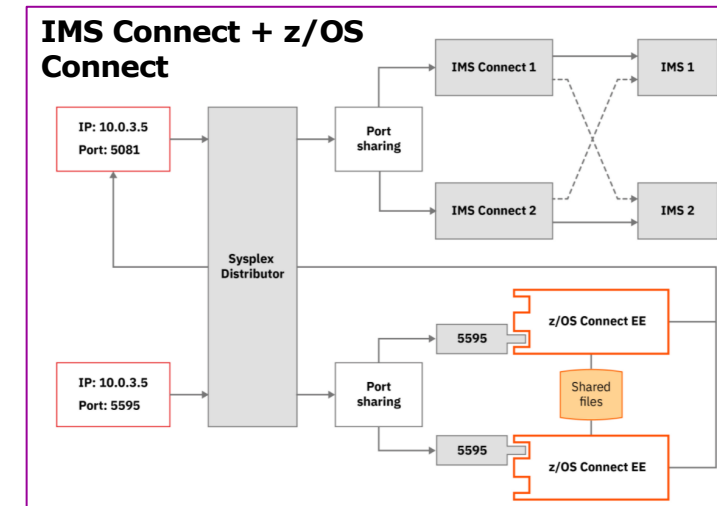
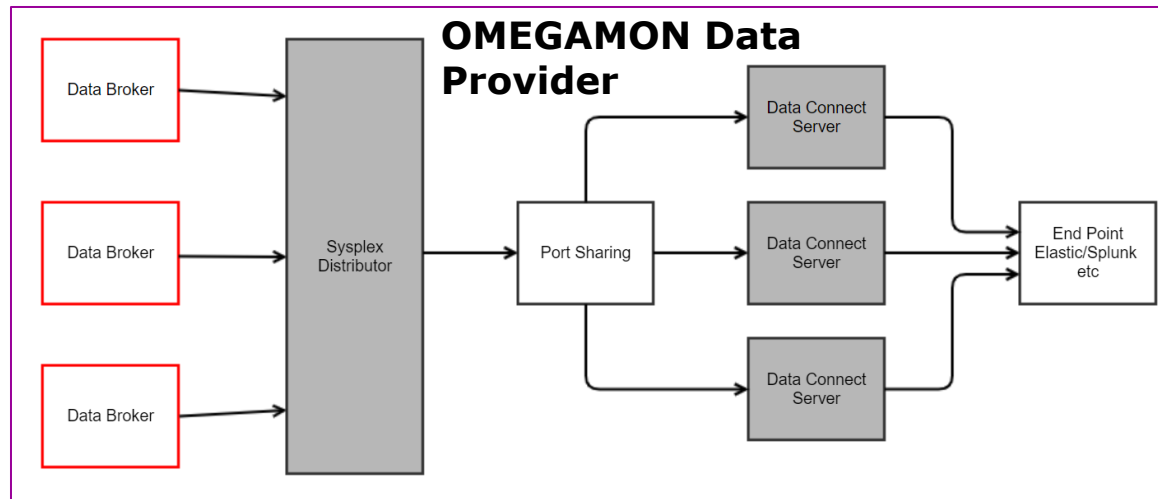
The architecture of OMEGAMON Data Provider allows for an expandable technology stack to scale implementations for added reliability and resilience.

Testing has proven technologies like Kafka can be clustered and incorporated into the stack to:

- Protect the OMEGAMON Data Connect Server from back-pressure from analytics engine data ingestion.
- Limit data loss due to component outage or failure.
- Provides scale for greater volumes of data.

**On Z,** Sysplex Distributor and Port sharing can be used to build redundancy into OMEGAMON Data Provider using proven techniques such as that used by z/OS Connect EE and IMS Connect.

**Off Z,** similar technologies and solutions exist.



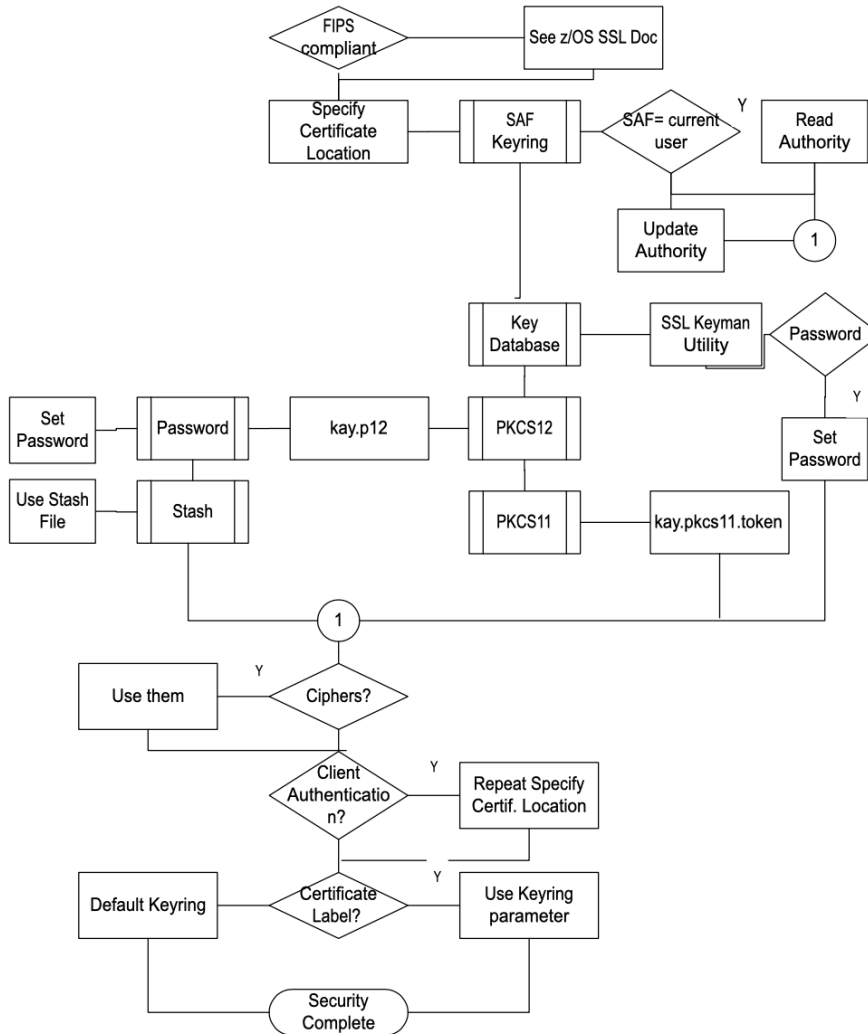
# Network Configuration For Security

---

The topology for deployment of the ODP infrastructure can vary as much as the deployment of analytic servers. As a result, there is a wide range of choices for Network security.

It's important to determine what you are protecting from:

# For each network connection, determine security flow



- Determine if FIPS compliant – this will define cipher type
- Where is the security certificate stored?
  - SAF keyring?
    - Identify what authority is necessary for key ring
    - Identify the keyring to be used
  - Key Database: leverage GSKKMAN utility
  - PKCS12: using a password or Stash file
  - PKCS11: token name
- If this is a Client authentication, then make the proper credentials for that request repeating above Security certificate process

# Data Broker client to Data Connect Server security

## Where does it fit

## Configure each side of connection


## TIPS

```
* SSL parameters
KAY.CIDB.FWD.OM.SECURITY=TLSv1.2
KAY.CIDB.FWD.OM.FIPS=ON|OFF
KAY.CIDB.FWD.OM.KEYRING=<string>
KAY.CIDB.FWD.OM.STASH=<string>
KAY.CIDB.FWD.OM.PASSWORD=<string>
KAY.CIDB.FWD.OM.CIPHERS=<string>
KAY.CIDB.FWD.OM.CERTLABEL=<string>
```

```
connect:
  input:
    tcp:
      enabled: true
      hostname: <connect_host>
      port: <connect_port>

enabled: boolean
ciphers: ciphers_list
client-auth: need|none|want
enabled-protocols: protocols_list
protocol: protocol
key-alias: string
key-password: string
key-store: string
key-store-password: string
key-store-type: JKS|PKCS12|JCERACFKS
trust-store: string
trust-store-password: string
trust-store-type: JKS|PKCS12|JCERACFKS
```

PARMLIB(ZWESIPxx)connect.yaml



1. PARMLIB for ODP Data Broker client
2. YAML for ODP Data Connector server

- Make sure that end-to-end workflow operates properly before adding network security parameters. This will minimize debug should an error occur.
- Ensure topology choices match at each end of network

# Data Connect Server to Analytics server security

e.g. Elastic, Splunk, Kafka


## Where does it fit

## Configure each side of connection

## TIPS

```
output:
  tcp:
    enabled: true
    sinks:
      logstash:
        hostname: <logstash_host>
        port: <logstash_port>
    enabled: boolean
    ciphers: ciphers_list
    client-auth: need|none|want
    enabled-protocols: protocols_list
    protocol: protocol
    key-alias: string
    key-password: string
    key-store: string
    key-store-password: string
    key-store-type: JKS|PKCS12|JCERACFKS
    trust-store: string
    trust-store-password: string
    trust-store-type: JKS|PKCS12|JCERACFKS
```

connect.yaml



1. YAML for ODP Data Connector Client connection to Analytics server
2. See Analytics server documentation for corresponding connection setup

- Make sure that end-to-end workflow operates properly before adding network security parameters. This will minimize debug should an error occur.
- Ensure topology choices match at each end of network
- Output type can be tcp: (Elastic, Splunk) or kafka:
- This is considered a PUSH style connection to server

# Analytics client to Data Connect server security

*e.g. Prometheus*

## Where does it fit

```
server:
  address: 0.0.0.0
  port: 9080
  ssl:
    enabled: true
    enabled-protocols: TLSv1.2
    protocol: TLS
    client-auth: need
    # Server certificate
    key-store: safkeyring:///STCOMDP/OMDPPring
    key-store-type: JCERACFKS
    key-store-password: password # Required fixed value
    key-alias: OMDPcert
    # Trusted client certificates
    trust-store: safkeyring:///STCOMDP/OMDPPring
    trust-store-type: JCERACFKS
    trust-store-password: password # Required fixed value
```


**connect.yaml**

## Configure each side of connection

1. YAML for ODP Data Connect Server connection to Analytics client
2. See Analytics client (e.g. Prometheus) documentation for corresponding connection setup

## TIPS

- Make sure that end-to-end workflow operates properly before adding network security parameters. This will minimize debug should an error occur.
- Ensure topology choices match at each end of network
- This is considered a PULL style connection to server





# Tips: Lessons learned

---

Zowe confusion – use an existing one – but forget to configure the ODP Data Broker plugin

Elastic set up – used a left over, complex logstash config instead of 18-line version

JZOS dependency – not properly installed or configured for ODP Data Connect server to operate

Misconfiguration of a non-z ODP Data Connect server instance(s)

Forget to turn on Near Term History collection within TOM or TEP prior to ODP operation

Use of an incorrect code page when editing YAML or proclib containing yaml input

STEPLIB for broker must be PDSE

Network configuration confusion – try it (PoC) without network security configuration to ensure it works before you add network security

Do not share a Docker volume across different Analytic server instances

Remove old Docker container instances when new ones become available

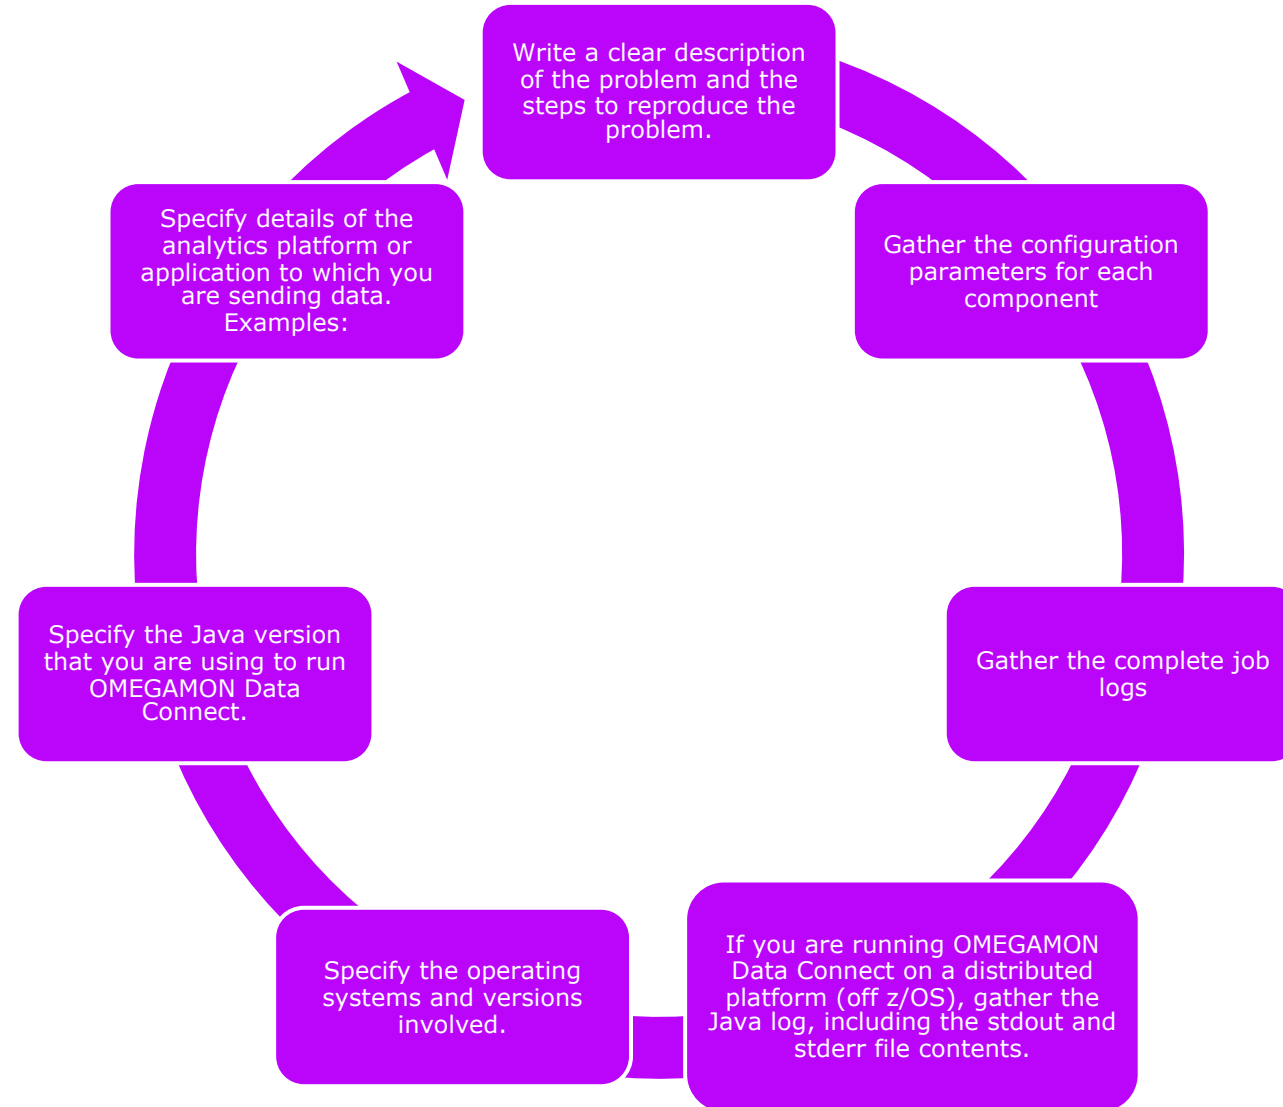
# Troubleshooting

# Basic Troubleshooting Tips

- If possible, before introducing SSL/TLS (security protocols), test that your configuration works without SSL/TLS. For example, in a sandbox environment that is entirely inside a secure intranet.
- Check that you are using the correct character encoding for each configuration member. For details, see Chapter 3, “Overview of configurable parts”.
- As a rudimentary test that OMEGAMON Data Connect is receiving the expected data from OMEGAMON Data Broker, temporarily enable the STDOUT output of OMEGAMON Data Connect.



# Gathering Diagnostic Data



# Gathering Diagnostic Data

1. Write a clear description of the problem and the steps to reproduce the problem.
2. Gather the configuration parameters for each component
  - **RKANPARU(KAYOPEN)** Collection configuration
  - **PARMLIB(ZWESISxx)** Zowe cross-memory server configuration, containing OMEGAMON Data Broker configuration parameters
  - **config/connect.yaml** OMEGAMON Data Connect configuration
3. Gather the complete job logs
  - The address spaces where the OMEGAMON collection tasks are running. For example, for the z/OS monitoring agent: the z/OS monitoring server address space.
  - The Zowe cross-memory server that is running OMEGAMON Data Broker.
  - OMEGAMON Data Connect, if you are running it on z/OS. Store each job log in a separate text file with a semantic (meaningful, plain English) name that identifies its contents (for example, include in the file names the terms "collection", "broker", "connect").
  - **Tip:** In z/OS SDSF, to save the complete job log to a data set, enter the action XD next to the job.
4. If you are running OMEGAMON Data Connect on a distributed platform (off z/OS), gather the Java log, including the stdout and stderr file contents.
5. Specify the operating systems and versions involved.
  - z/OS version
  - If you are running OMEGAMON Data Connect off z/OS, the corresponding details for that platform, such as the operating system distribution name and version.
6. Specify the Java version that you are using to run OMEGAMON Data Connect.
  - **Tip:** To get the Java version, use the command **java -version**.
7. Specify details of the analytics platform or application to which you are sending data. Examples:
  - The name and version of the analytics platform.
  - The operating system distribution name and version.
  - How you have configured the analytics platform to ingest data from OMEGAMON Data Connect. For example, for the Elastic Stack: the Logstash configuration and index template; for Splunk, the configuration stanzas.
  - Whether, and how, you have tested that the destination is correctly configured to ingest data, independent from OMEGAMON Data Provider. For example, have you used a stand-alone TCP forwarder to send a sample line of JSON to the destination, in the same format sent by OMEGAMON Data Connect



# Thank you.

[rocketsoftware.com](https://rocketsoftware.com)

[zconcierge@rocketsoftware.com](mailto:zconcierge@rocketsoftware.com)



© Rocket Software, Inc. or its affiliates 1990 – 2023. All rights reserved. Rocket and the Rocket Software logos are registered trademarks of Rocket Software, Inc. Other product and service names might be trademarks of Rocket Software or its affiliates.  
© Copyright IBM Corporation 2023. IBM, the IBM logo, ibm.com, and Watson are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).