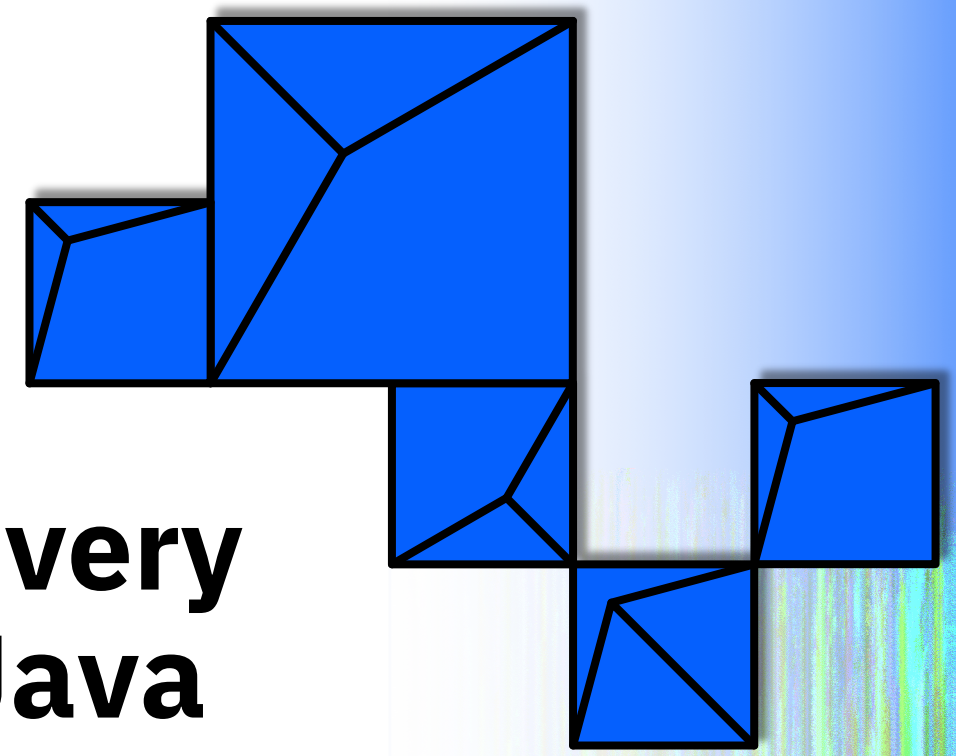

z/TPF Guaranteed Delivery for Java Services and Java Logging enhancements

Daniel Gritter

Disclaimer

Any reference to future plans are for planning purposes only. IBM reserves the right to change those plans at its discretion. Any reliance on such a disclosure is solely at your own risk. IBM makes no commitment to provide additional information in the future.

Agenda

Background

Problem Statement

Pain Points

Value Statement

Technical Details

Conclusion

JAM Support

REST service support for Java, using JAX-RS interface (PJ43892 – delivered 1Q2017)

- Supports internal and external REST clients
- tpf_srvcInvoke
- Request / Reply model
- Synchronous support only

Problem Statement

As a **z/TPF architect**, I want to **publish data** from z/TPF through a Java environment with an **at-least-once delivery** guarantee, without sacrificing throughput or introducing bottlenecks.

Pain Points

- tpf_srvcInvoke calling a Java service currently has two options:
  - Return immediately (can’t guarantee delivery)
  - Block until delivery (affects throughput / existence time)
- No out-of-the box solutions
- Performance impacts of REST service overhead

Value Statement

An enterprise architect can seamlessly integrate and guarantee delivery of TPF data across their enterprise using an industry-standard solution in under a week

Technical Details

- Guaranteed Delivery for JAM provides:
  - MQ Connector to Java services
  - IBM Commit Scope support to protect in-flight data
  - Custom and IBM-provided solutions

- IBM provided Kafka connector (PJ45923, delivered March 2020)
- IBM provided SMTP connector (PJ46000, delivered May 2020)

Technical Details

- TPF Application writes formatted Guaranteed Delivery for JAM (GD4J) request to local queue.
  - Topic Name, Key Data, Body Data
- Worker thread from one JVM in one JAM reads the message under commit scope.
- Only one JAM should process a given queue.
- Publish only – application does not receive a response

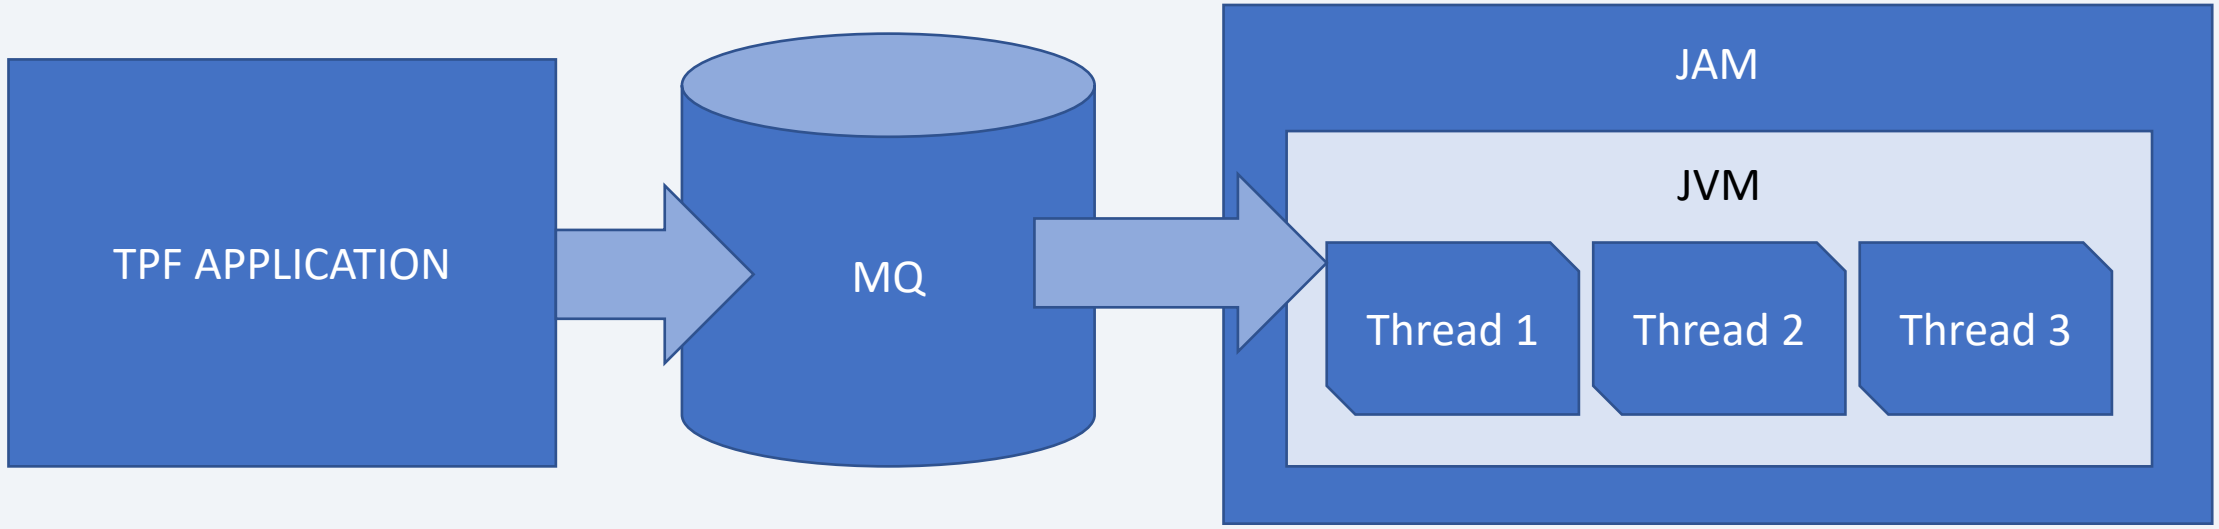

Technical Details

- GD4J Router reads from MQ and dispatches to connector class
- IBM provided connector routes to Kafka topic
  - Topic name defines connectivity information
- IBM provided connector routes to SMTP server
- Customer written connector routes to custom Java code

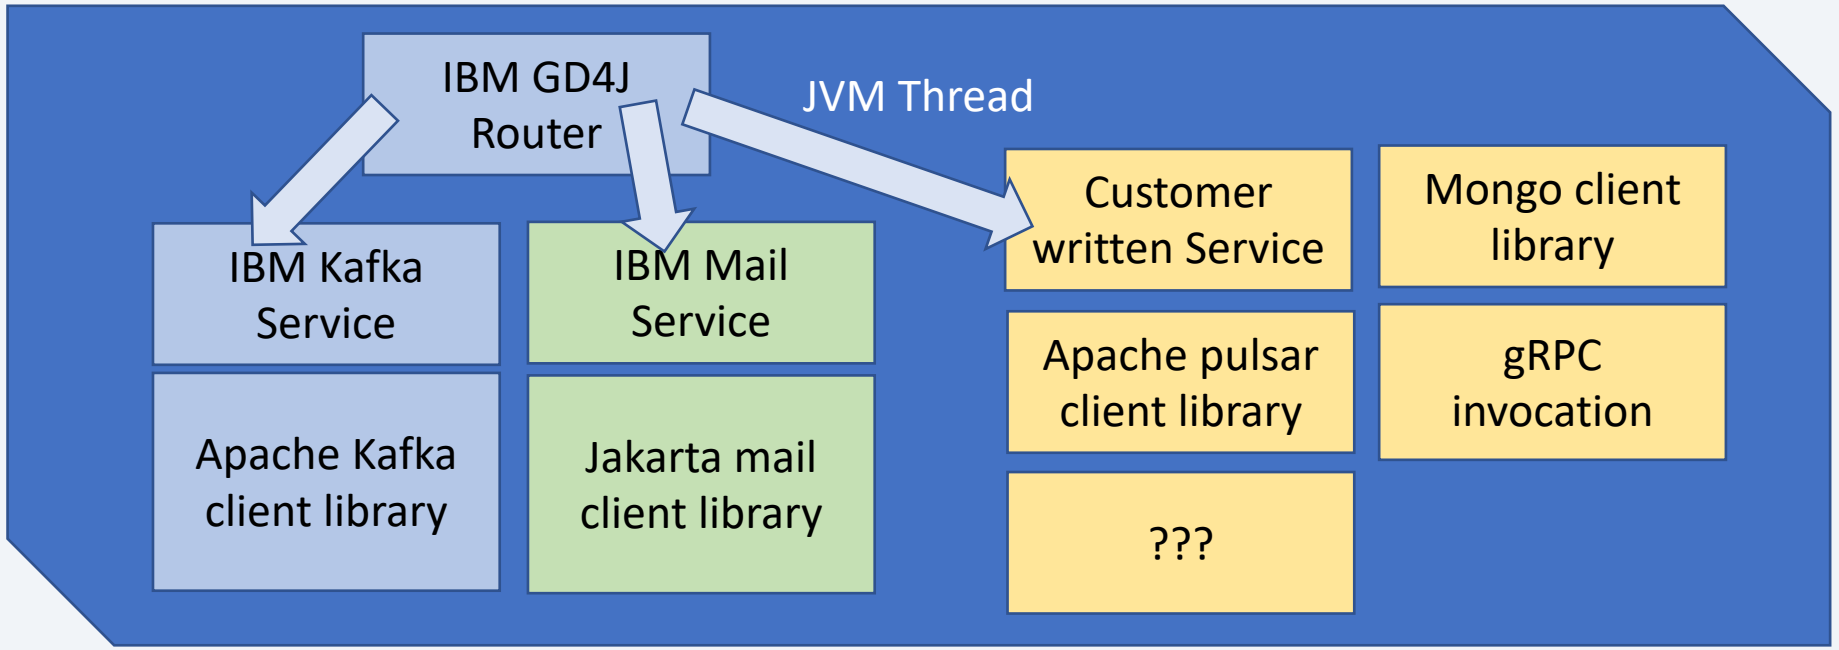

Technical Details

- Connector Interface
- Supports Synchronous and Asynchronous Java programming models
- Synchronous model maintains commit scope until return from connector.
- Asynchronous model detaches commit scope from worker thread so that it can complete asynchronously
- PublishRequest.async() transforms to asynchronous mode
- Must issue PublishRequest commit, rollback, retry, error or async api prior to return from publish method.

```
publish(PublishRequest req) {
        … do some work
        req.commit();
}
```

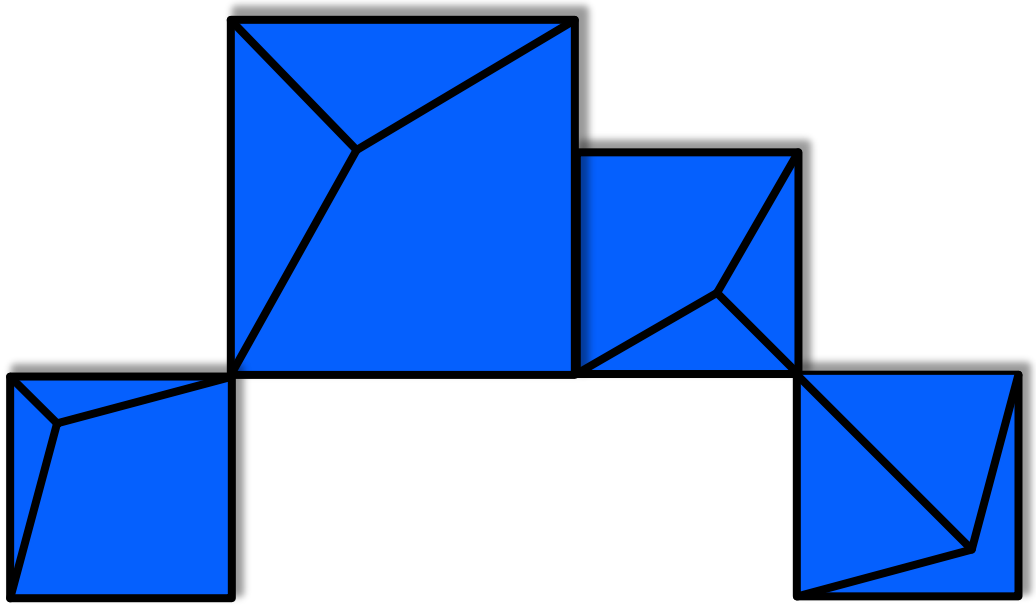

Logging through log4j2

Problem Statement

Logging for a JAM only gets sent to the “stderr.log” and “stdout.log” for each JVM, but it’s difficult to track multiple files in the file system for updates

Technical Details

- PJ45923 introduces support for log4j2 as part of JAM processing.
- New log4j2.xml configures how output is routed.
- Uses /dev/tpf.omsg (WTOPC filesystem driver) to route messages to the TPF console
  - By default, will only print one message per minute
- Full logs still placed in stderr.log for each JVM
- Users can either edit the IBM-provided base/tpfjax/log4j2.xml or load unique configurations in a JAM using a filesystem classpath or JAR file including log4j2.xml as a resource

Technical Details

Log4J2 configuration

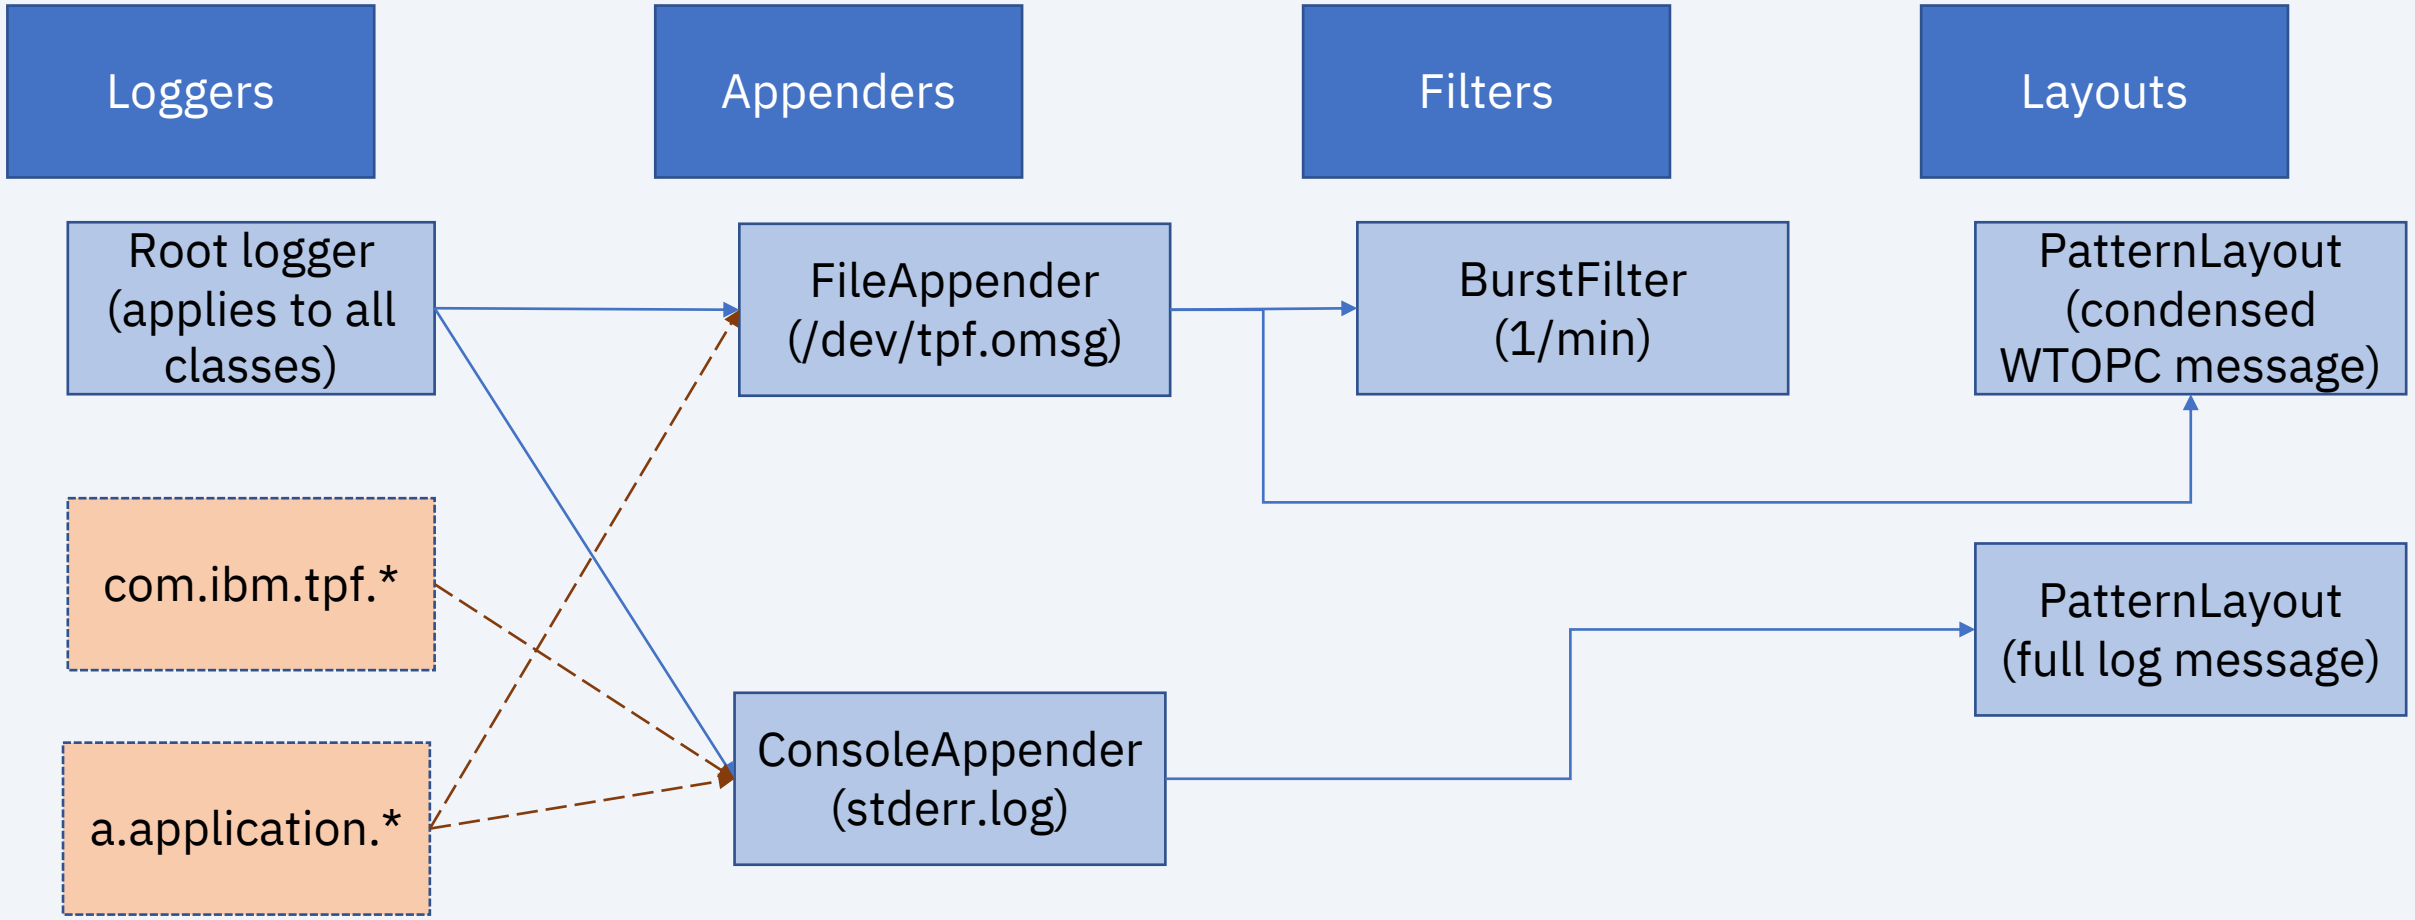

Conclusion

- Using the Guaranteed Delivery for JAM simplifies data publishing and other enterprise integration methods
- MQ allows use of IBM Commit Scope to guarantee delivery through Java without impacting application throughput
- IBM Provided connector allows for integration with Kafka + SMTP
- User-written connector allows for custom integrations (MongoDB, Apache Pulsar, gRPC, and more)
- In-flight limit allows for smoother remote system outage / failover handling and throttling
- Log4j2 allows customers to configure how messages from Java get logged
- PJ45923 available March 2020

Thank You

Questions? Comments?

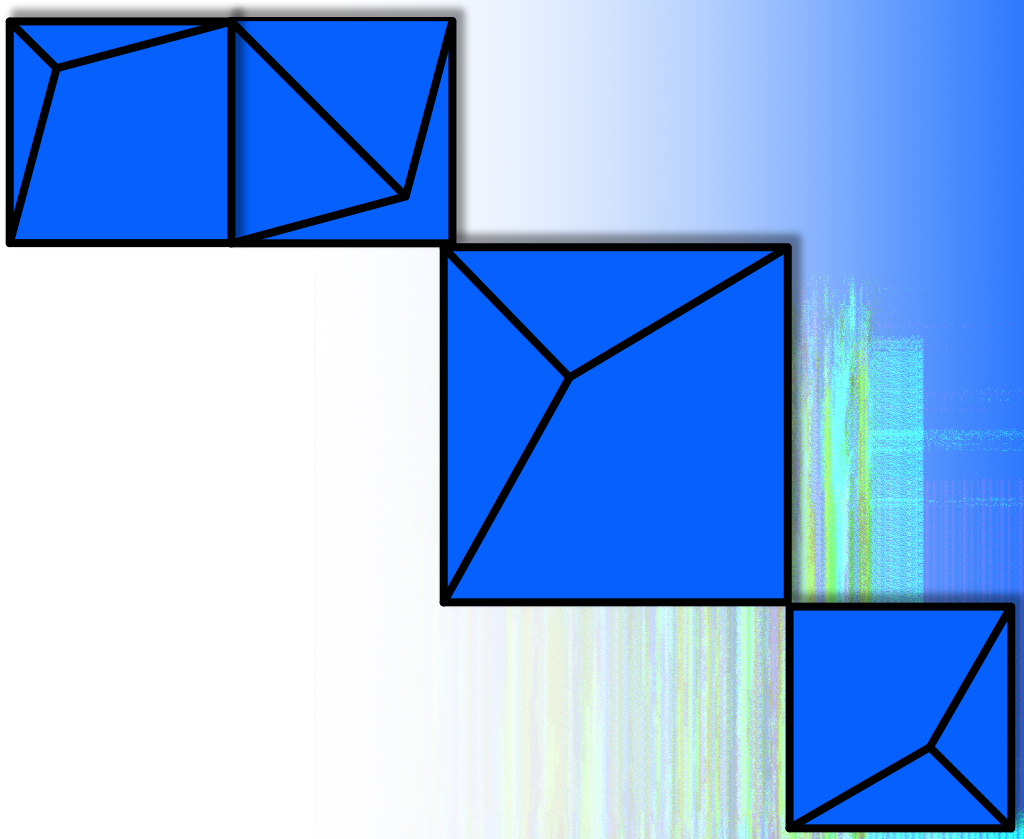

Virtual TPFUG Q&A

Summary of Q&A from the virtual TPFUG event:

Question	Answer
Q1: So Java monitoring will use Java-based Kafka support on z/TPF? Q2: So the Kafka component is on the z/Linux side? (slide 17)	A1: JVM monitoring support will not use the Java-based Kafka support delivered in PJ45923 that resides on z/TPF; it will use the RTMC infrastructure. The RTMC infrastructure will send the data off of z/TPF to a Kafka server on z/Linux, however. A2: Correct, the data goes from Java (z/TPF) --> RTMC (z/TPF) --> Kafka (z/Linux).
Q: Does the 10 sec JAM startup time include the time to reload the cache, or does the timer start when you IPL the JVMs and assumes the cache is available? (slide 22)	A: The 10 seconds assumes the cache was recovered. This allows the shared cache to already be available across a system IPL and the JVMs simply need to access it again, but not rebuild it.

Trademarks

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon 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 given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-