Runtime Metrics Collection Update

2023 TPF Users Group Conference
April 24-26, Dallas, Texas
Operations and Coverage

—

Josh Wisniewski

IBM

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.

What is runtime metrics collection?

Runtime metrics collection facilitates the following monitoring levels:

- System – Is the overall z/TPF system healthy? CPU usage, block usage, etc.
- Application – The system can be healthy, but the applications could be rejecting all transactions or a given message type using too much CPU.
- Workload (future) – Provides metrics for application performance monitors to more easily diagnose workload issues in the enterprise context.
- Business – Health of the business and maximizing results. Is a sale working? Add inventory. Is a fare underperforming? Adjust the fare.
- Enterprise – z/TPF is part of the bigger picture.

What is runtime metrics collection?

Runtime metrics collection provides various capabilities for the different monitoring levels:

- System – CDC, JVM, user-define metrics* (like # proprietary API calls).
- Application – Name-value pair collection (msgtype, code package, etc), JVM, user-define metrics* (like application error rates, timeouts, etc).
- Workload – Future.
- Business – User-define metrics* (sales, inventory, etc. metrics).
- Enterprise – Subset of CDC, user-define metrics*, etc.

* User-defined metrics release target 2Q2023.

What is runtime metrics collection?

- Provides continuous real-time monitoring and analysis of a z/TPF system.
- Provides system wide monitoring of Java on z/TPF.
- Collection leverages sampling so it can run continuously without impacting the z/TPF system.
- z/TPF real-time insights dashboard starter kit provides sample analytics pipeline with dashboards. https://www.ibm.com/support/pages/ztpf-real-time-insights-dashboard-starter-kit
- You can implement machine learning and data science techniques to provide more insights into resource usage.
- Message analysis tool results analysis and dashboards (PJ46308). Provides resource usage insights at the function and macro level.

Delivered since the last TPFUG:

- PJ46737 (June 2022) – Complex-wide dashboard support
  - Metrics from multiple processors on a single dashboard including averages and totals across processors.
  - Includes a lot of valuable features for single processor environments.
    - New dashboards (system state, name-value pair, REST, business events).
    - New metrics (low priority CPU, workload CPU, etc).
    - MySQL database support (alternative to MariaDB).

Delivered since the last TPFUG:

- PJ46739 (Sep 2022) – Multiprocessor support for CDC collection type
  - Up to 10 processors can send CDC data to a single tpfrtmc offline utility instance.

Notes:

- Up to 10 processors can simultaneously send CDC (PJ46739), JVM (PJ46608), and message analysis tool (PJ46308) data to a single tpfrtmc offline utility instance.
- A single tpfrtmc offline utility instance is limited to processing name-value pair data from a single z/TPF processor. Multiprocessor support for the NVPC collection type is targeted for 2H2023.

Delivered since the last TPFUG: 5 LPARs, actual CPU, avg

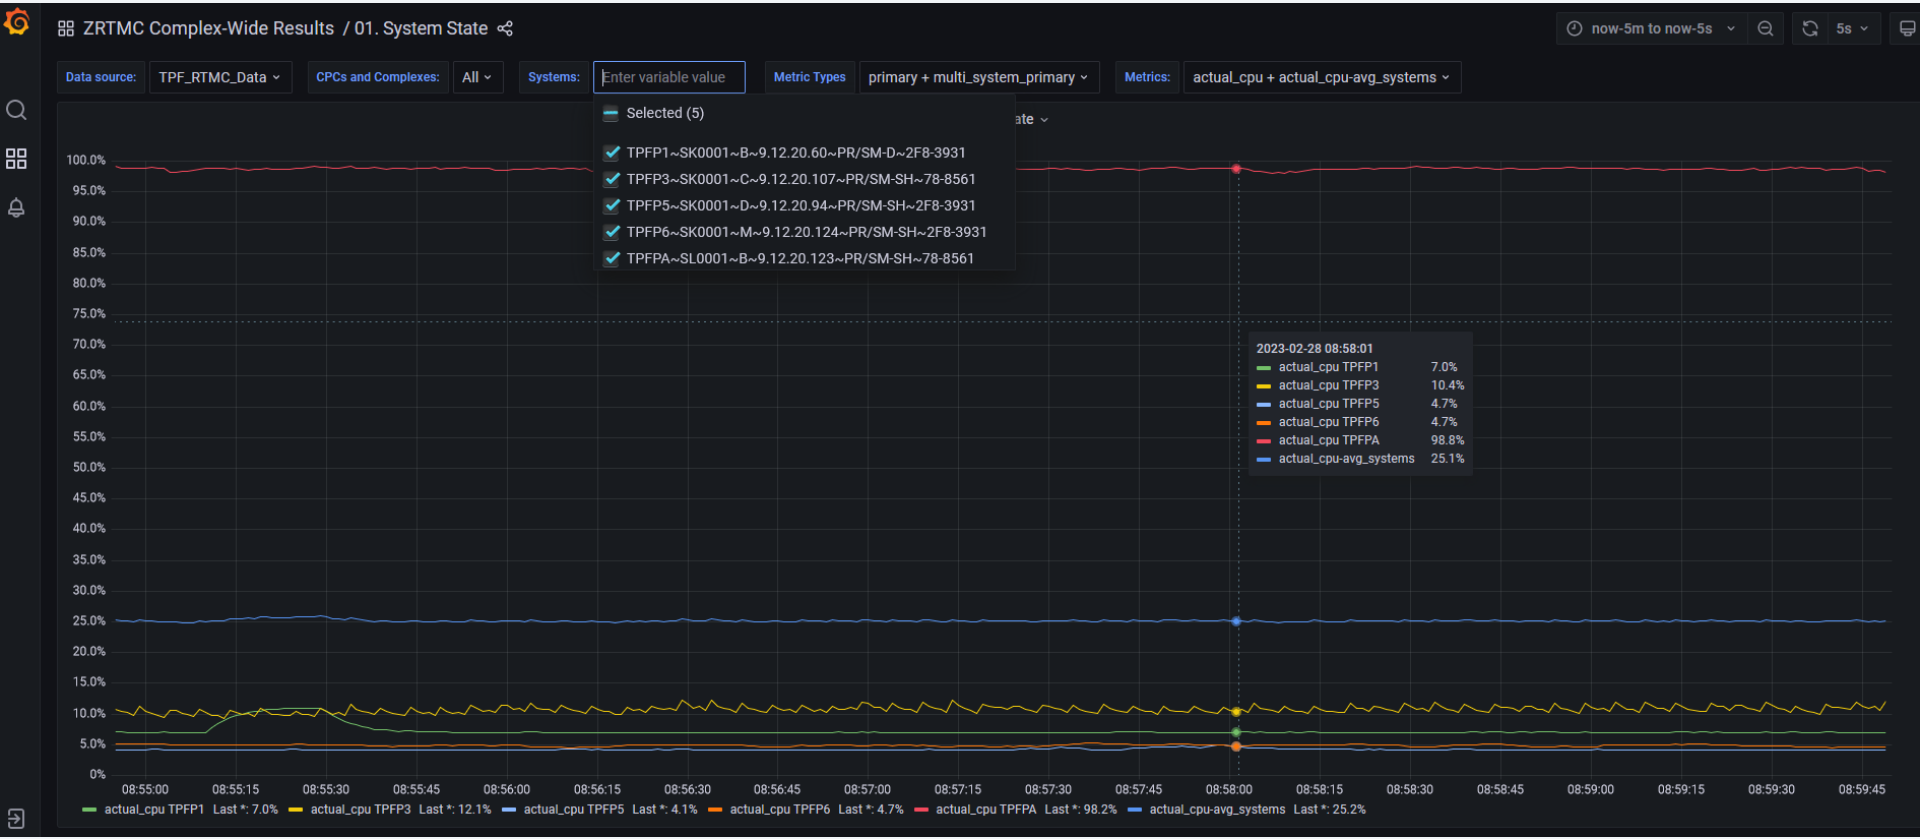

Delivered since the last TPFUG: business events

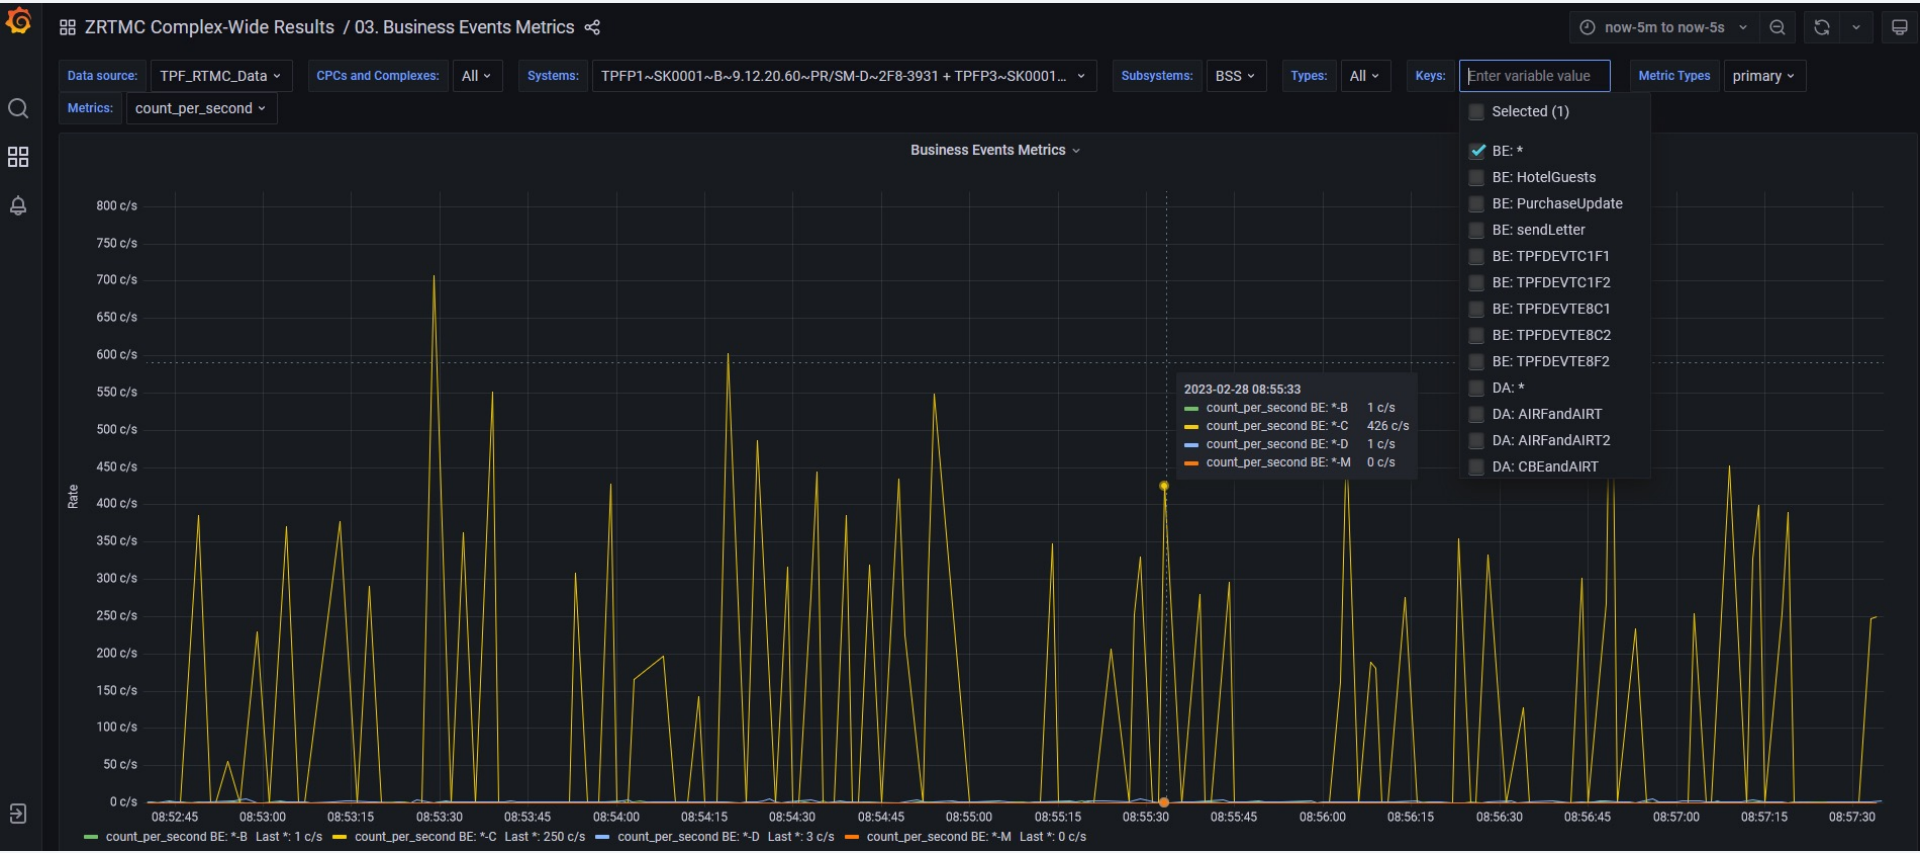

Delivered since the last TPFUG: REST Services

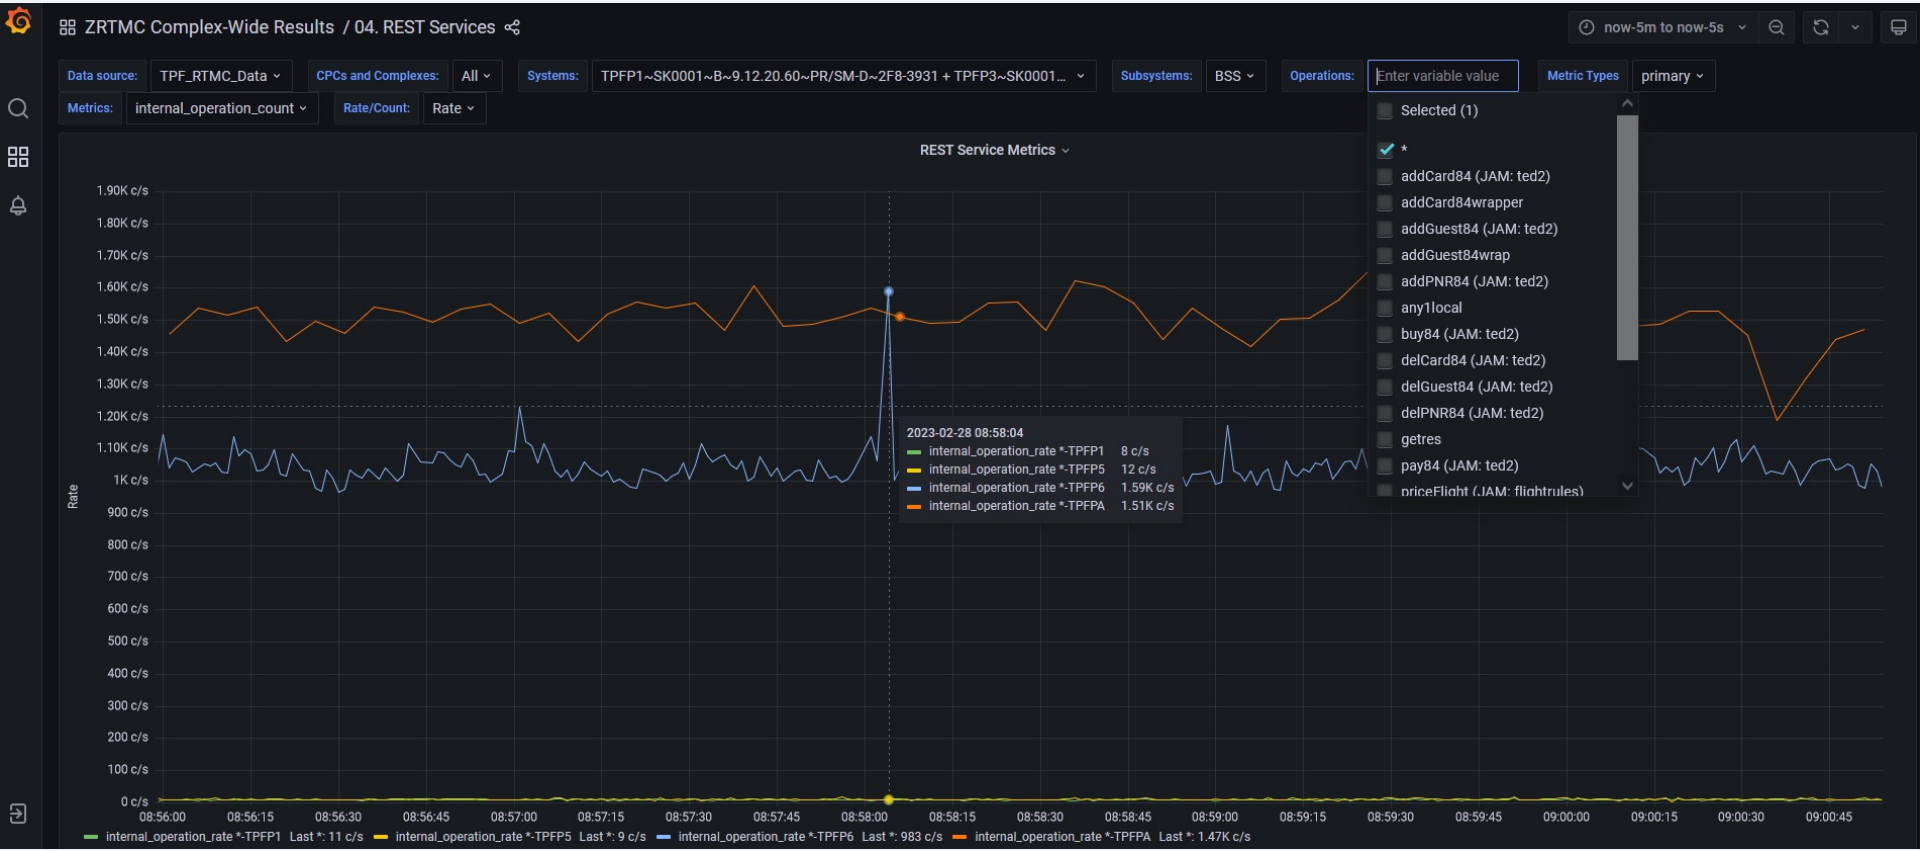

Delivered since the last TPFUG:

- PJ47008 (Mar 2023) – Secure Kafka network connections
  - In the normal operation of runtime metrics collection, the tpfrtmc offline utility sends data to Apache Kafka and the tpf_zrtmc_analyzer reads data from Apache Kafka. Some customers require these connections to Apache Kafka be encrypted in production environments.

Delivered since the last TPFUG: Secure Kafka network connections

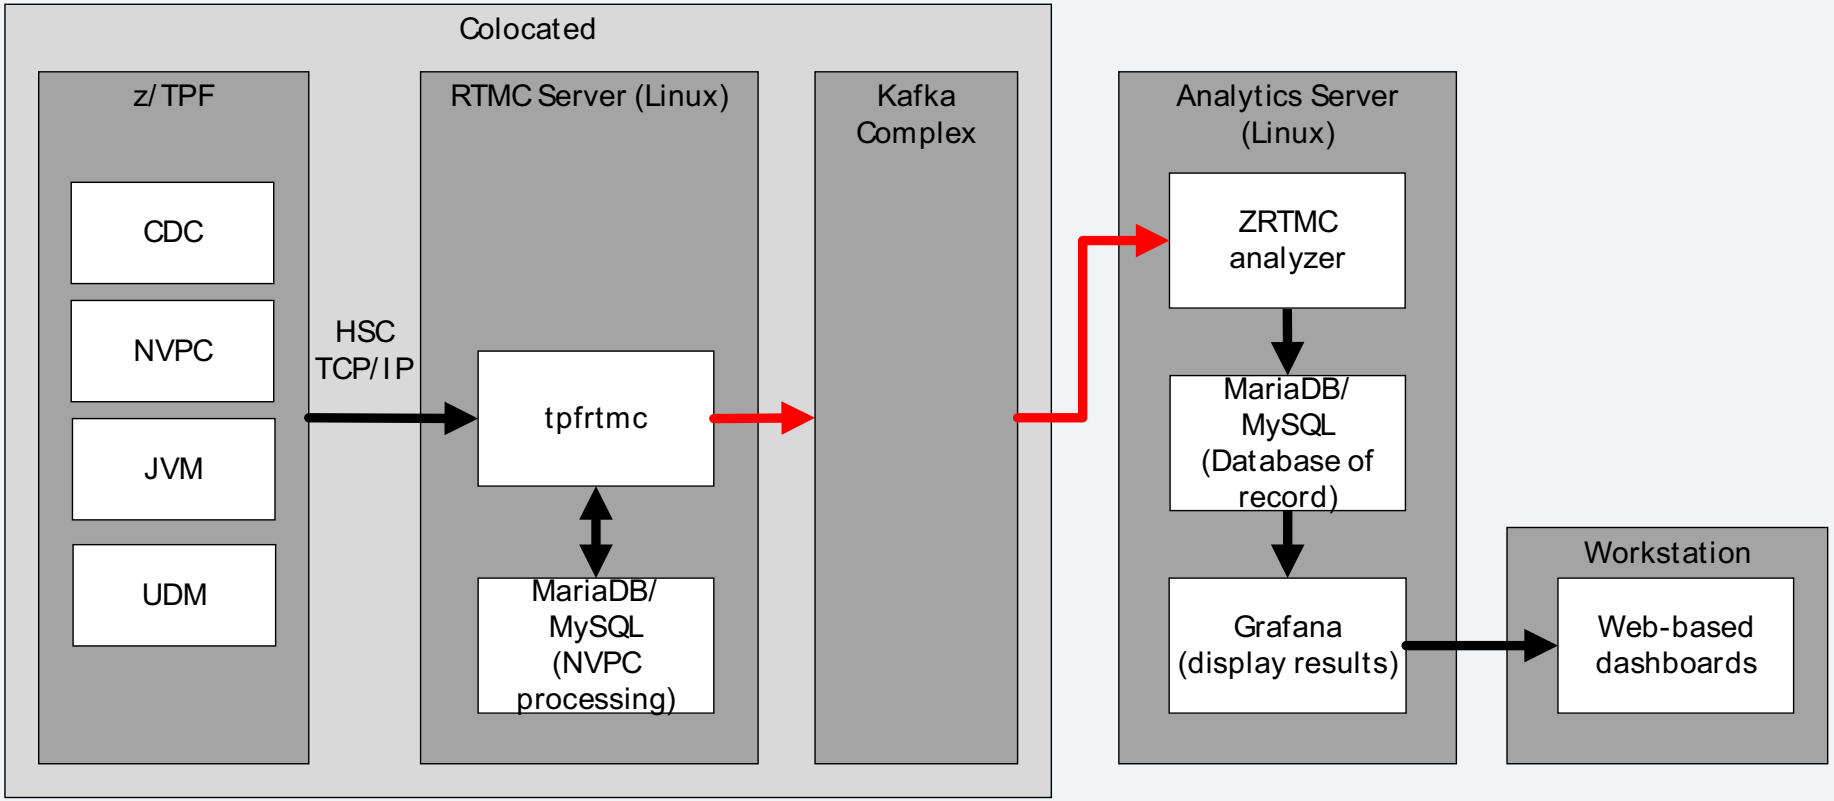

Delivered since the last TPFUG:

A primary focus has been facilitating customer adoption into production environments. As part of these efforts, the following APARs are of interest:

- PJ46876 (Oct 2022) – Identified MySQL server memory leak processing name-value pair data and provided installation instructions to prevent the issue.
- PJ46853 (Nov 2022) – Integer name-value pair skip factor changed to float scale factor for better precision in various metric calculations.
- PJ46898 (Sep 2022) – Fixed malformed data issue with greater than 49 I-stream processor configurations.
- PJ46946 (Dec 2022) – Reduce required database permissions for MariaDB and MySQL databases.

2023 Roadmap:

Our top priority will continue to be facilitating customer adoption and moves into production environments. In addition to answering questions, we've addressed several blockers to adoption identified by customers:

- MySQL support.
- Secure Kafka network connection.
- Reduce database permissions required.

We are currently working on resolving these customer identified blockers:

- Linux on IBM Z Docker configuration.
- Java 11
- Secure network connection from z/TPF to offline components.

2023 Roadmap:

Target 2Q2023 – Linux on IBM Z Docker configuration

- Leverages IBM Z Container Registry for trusted open-source container builds.
- Additional features for Linux on IBM Z or x86 Linux:
  - Easier to configure installation of open-source dependencies from local or trusted sources: Docker registry, pypi, apt-get, apk.
  - Java 11 support.
  - All name-value pair metrics in name-value pair dashboards.
  - Sample analytics server database pruning to manage hard drive space.

2023 Roadmap:

Target 2Q2023 – User-defined metrics

- You can define and send operational, application, and business metrics for real-time monitoring, analysis, and data science.
  - Modify your online monitors to convert the data into JSON format.
  - Call the new API to send the metrics to the sample analytics pipeline: tpf_rtmc_send_user_defined_metrics
  - Implement your custom analysis and dashboards in the sample analytics pipeline.

2023 Roadmap: User-defined metrics sample dashboard

2023 Roadmap:

Target 2Q2023 – User-defined metrics

- 10 processors can send user-defined metrics data simultaneously.
- Step by step tutorial included in the IBM documentation with sample code.
- Secure network connection between z/TPF and the tpfrtmc offline utility.

2023 Roadmap: User-defined metrics – Secure network connection

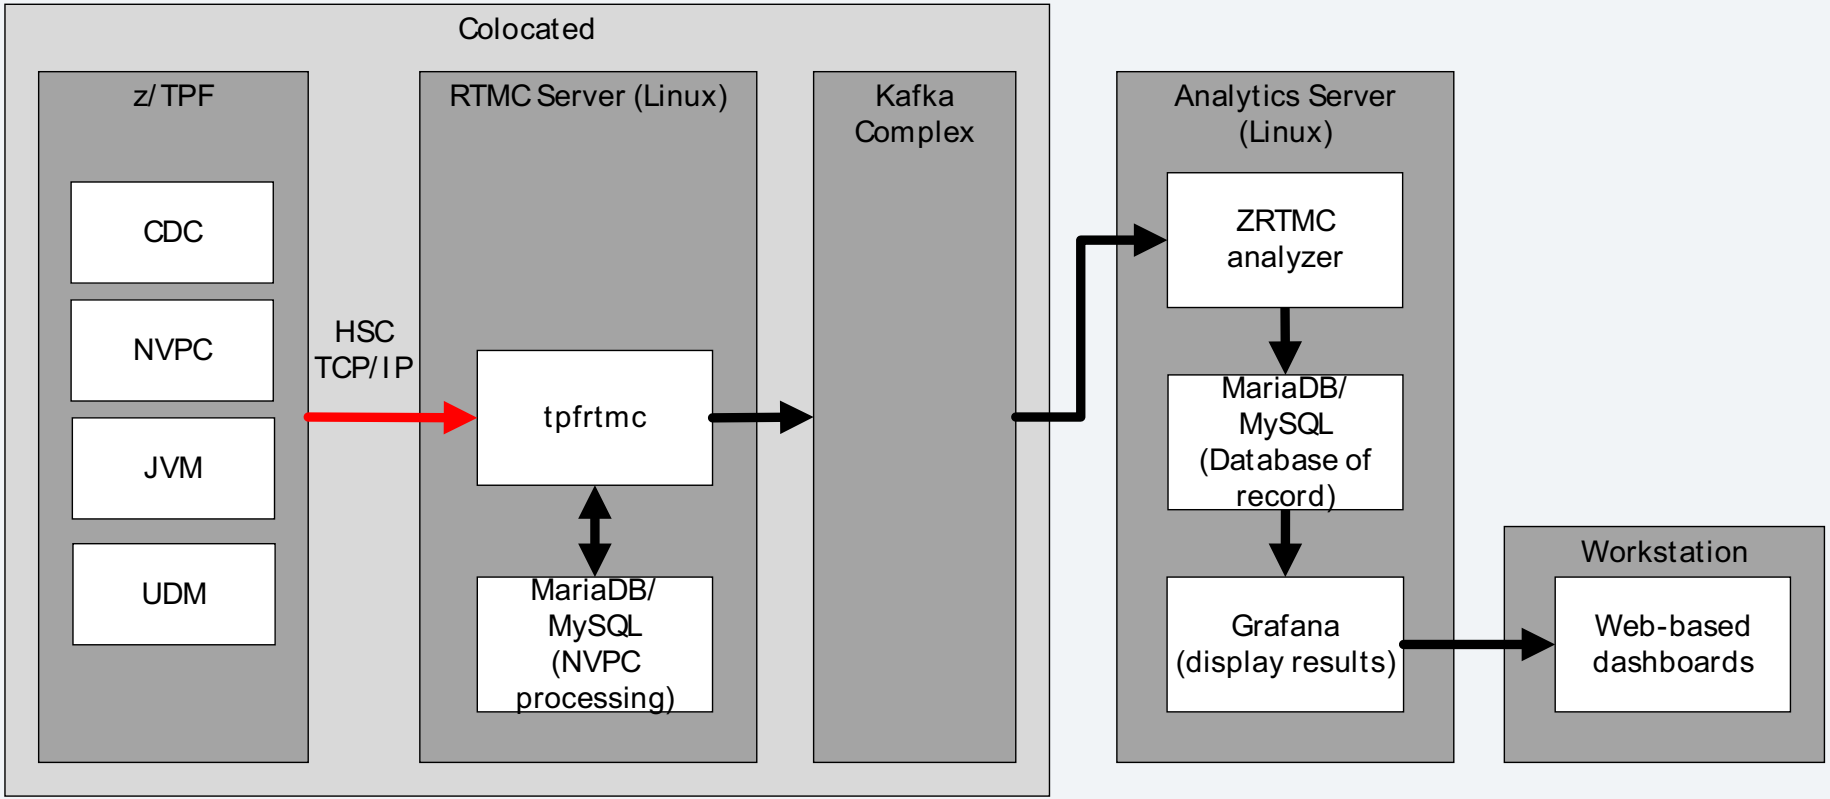

2023 Roadmap:

Target 2H2023 – Multiprocessor support for name-value pair collection

- Up to 10 z/TPF processors sending name-value pair collection data to a single tpfrtmc offline utility instance.
- Note: Up to 10 LPARs, 10 VPARs or any combination of sending all collection types (CDC, name-value pair collection, JVM, user-defined metrics, message analysis tool) simultaneously.

2023 Roadmap:

Target 3Q2023 – Improved high availability support

- Documented recommendation for high availability architecture.
- Improved behavior for various fail over scenarios.
- Starter kit replaced by fully supported sample analytics pipeline:
  - Delivered as part of the z/TPF code base.
  - Extensible just like the starter kit.
  - Will include improvements to name-value pair based metrics for low volume messages.
- ZRTMC can be run in CRAS state.
- This APAR must be loaded with the offline image loader (TLDR).

2023 Roadmap:

Ongoing in 2H2023 – Replace CDC Tivoli dashboards in Grafana

- Only two original CDC Tivoli data types are not supported by runtime metrics collection today: z/TPFDF and user-defined metrics.
- Several original CDC Tivoli data types do not have runtime metrics collection dashboards today.
- This effort will consist of several APARs delivering the missing features so that the IBM Tivoli Monitoring Agent for z/TPF support can eventually come to an end.
- This will include TE vs GP metrics at the system level on the complex-wide dashboards > system state dashboard.

2023 Roadmap: Target production architecture

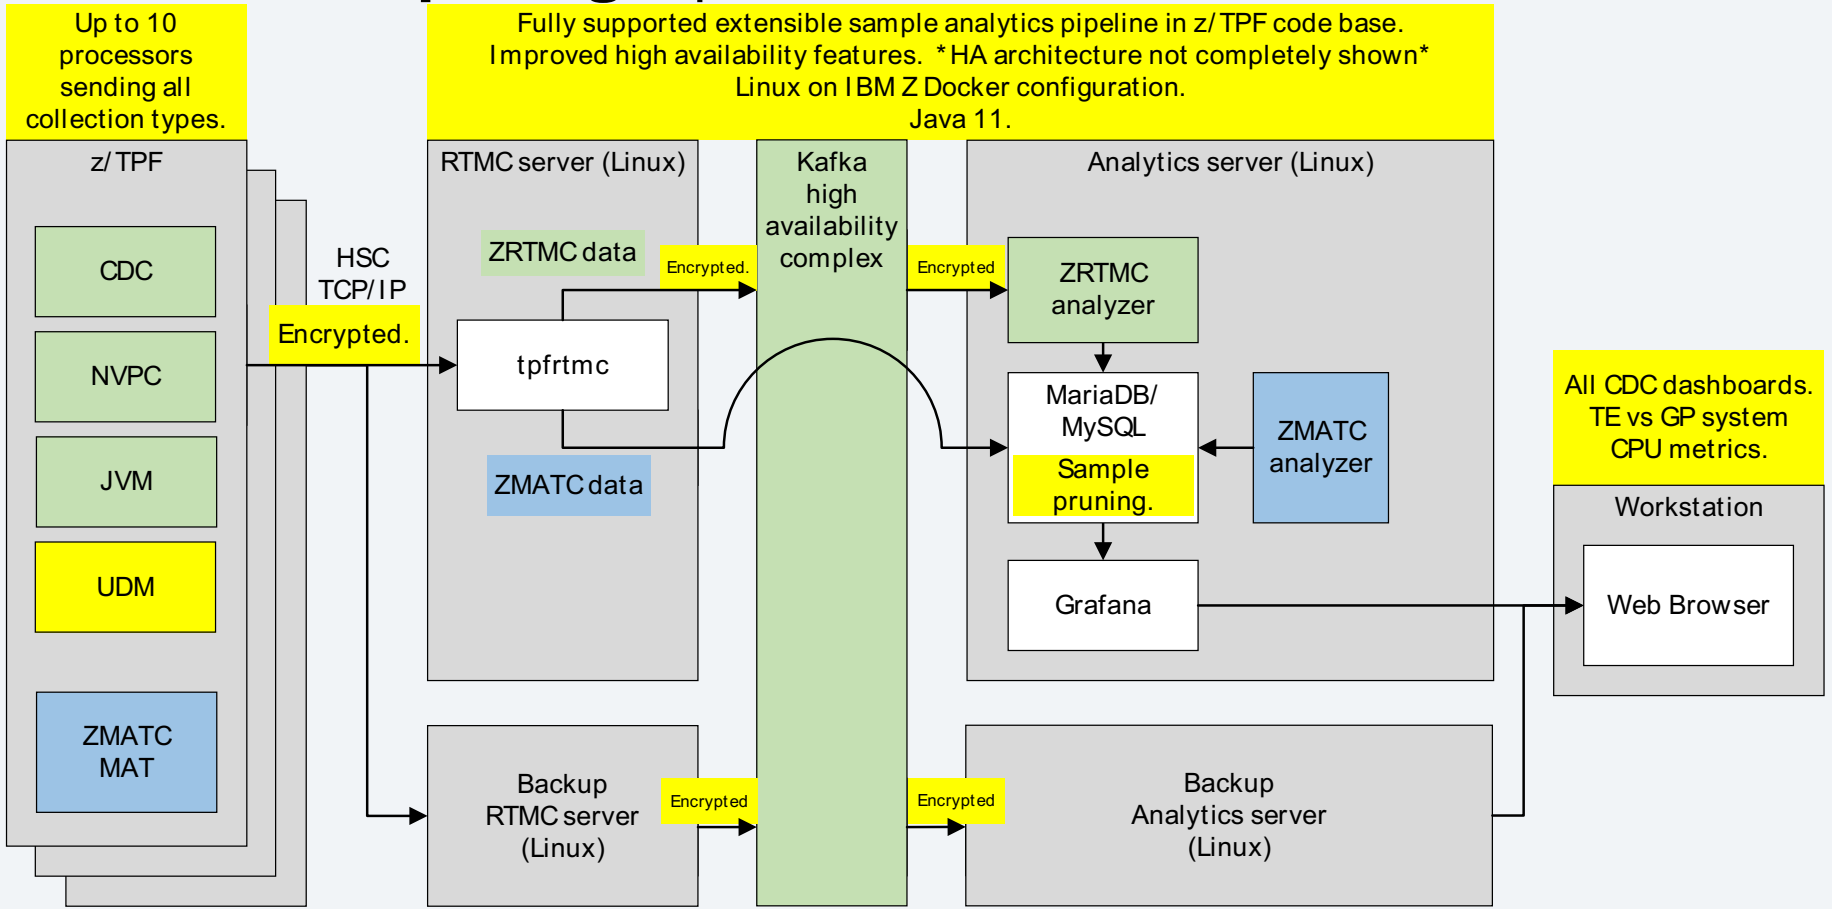

Feature list summary: Delivered

Enhancements:

PJ46737 (June 2022) – Complex-wide dashboard support, MySQL.

PJ46739 (Sep 2022) – Multiprocessor support for CDC collection type.

PJ47008 (Mar 2023) – Secure Kafka network connections.

Bug Fixes:

PJ46876 (Oct 2022) – MySQL server memory leak.

PJ46853 (Nov 2022) – Float name-value pair scale factor.

PJ46898 (Sep 2022) – Greater than 49 I-streams.

PJ46946 (Dec 2022) – Reduce required database permissions.

Feature list summary: 2023 Roadmap

Target 2Q2023 – Linux on IBM Z Docker configuration, Java 11, sample analytics database pruning.

Target 2Q2023 – User-defined metrics, secure z/TPF network connection.

Target 2H2023 – Up to 10 name-value pair processing.

Target 3Q2023 – Improved high availability support, fully supported extensible sample analytics pipeline.

Ongoing in 2H2023 – Replace CDC Tivoli dashboards in Grafana, TE vs GP metrics at the system level.

We want sponsor users!

Our development cycle is driven by your feedback.

We are looking for sponsor users to assist in design and implementation, targeting the following personas:

- Operations
- Coverage
- Offline Linux Administrators

We expect to have ongoing sponsor user engagements.

If you are interested in participating as a sponsor user, please contact:

Josh Wisniewski (jwisniew@us.ibm.com)

Thank you