


Real-time Runtime Metrics Collection and Real-time Insights Dashboard Starter kit



Josh Wisniewski

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

Name-value Pair Collection background
Real-time Runtime Metrics Collection
Real-time Insights Dashboard Starter Kit
Demo
Conclusion
Value Statement
What's next?

Background

- Name-value pair collection origin:
 - Data collection/reduction, Continuous Data Collection (CDC) and other tooling provides lots of insights into system wide resource usage.
 - However, insights into message and application resource usage was largely unavailable.
 - Customers requested a mechanism be provided to fill this gap.

Background

- Name-value pair collection provides the ability to collect system resource usage from many perspectives:
 - message type.
 - message origin.
 - code package.
 - and more, all defined by you!
 - And you can analyze the results from multiple perspectives at the same time.
- PJ44321 Oct/2017 and PJ45264 Jul/2018

Problem Statement

Name-value pair collection results are captured for a limited duration to tape which prohibits this type of collection from being used for real time monitoring.

Technical Details

- Real-time Runtime Metrics Collection
 - Data flows continuously in real time so that data is always immediately available (no waiting for collection, post processing or use of tapes).
 - Provides name-value pair collection and continuous data collection (CDC) data (that is, message/application level and system level data).
 - Name-value pair collection leverages sampling so it can run continuously without impacting the system.
 - Allows you to store historical data in a database to investigate situations that have occurred in the past.

Pain Points

- Analyzing data pain points:
 - The amount of diagnostic data is overwhelming when trying to solve problems thereby increasing the risk that investigations will take longer.
 - It often requires someone with decades of experience to solve problems quickly making it difficult to train up someone new.
 - Data is great, but what is needed is insights to speed up investigations.

Value Statement

Coverage programmers can leverage data science insights into system state and application trends in real time to prevent impacts to service level agreements.

Technical Details

- Real-time Insights Dashboard Starter Kit details:
 - Sample environment and analytics pipeline to play with and inspire your analysis. Provides statistical analysis.
 - View data and insights instantly on interactive dashboards.
 - Includes replay scripts so you can experiment without having to set up name-value pairs or install real-time runtime metrics collection on your z/TPF system.
 - Real NVP and CDC test data.
 - Component test data and usage instructions are included.
 - Leverages Docker on x86 Linux to make it easy to adopt, configure and use.
 - Docker setup for processing ZCNVP name-value pair collection results too.

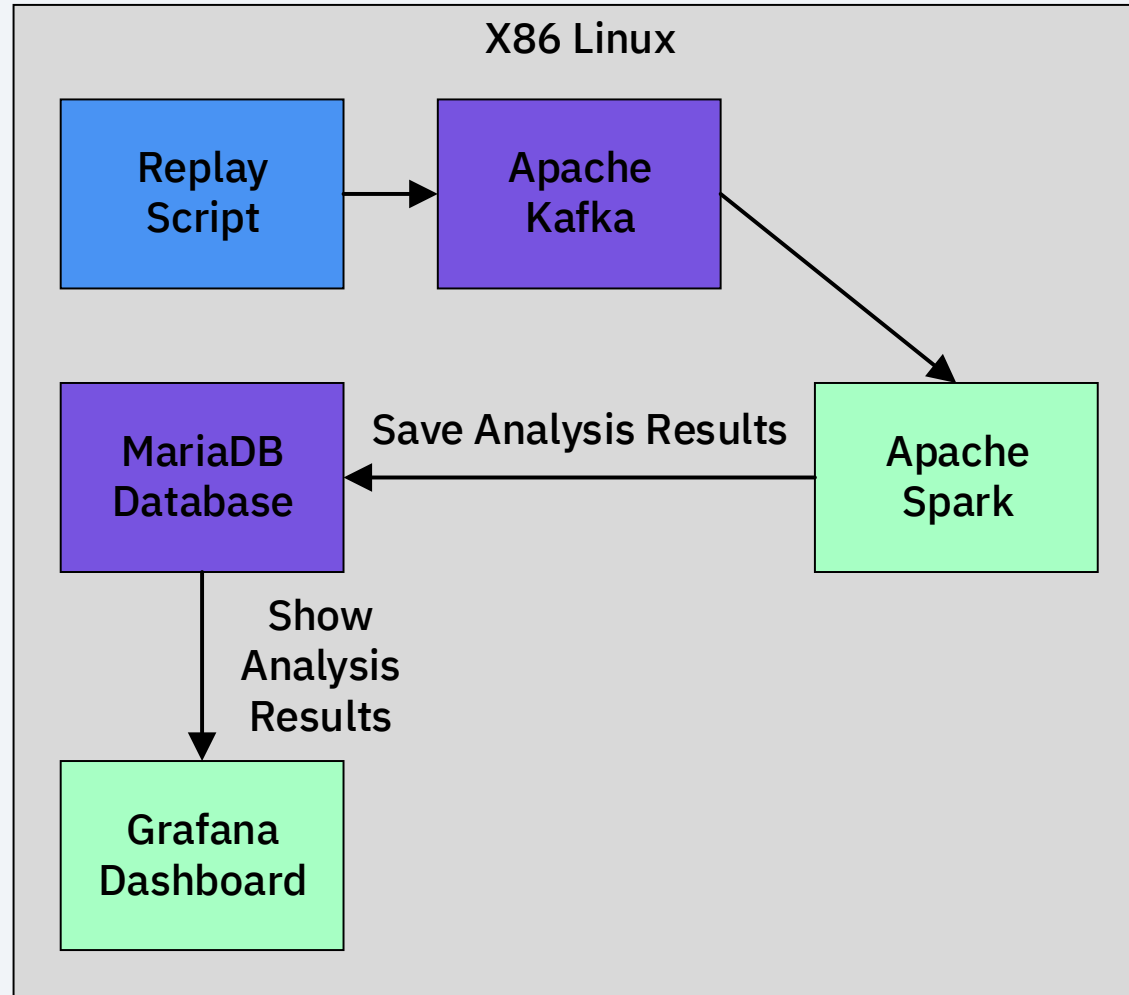
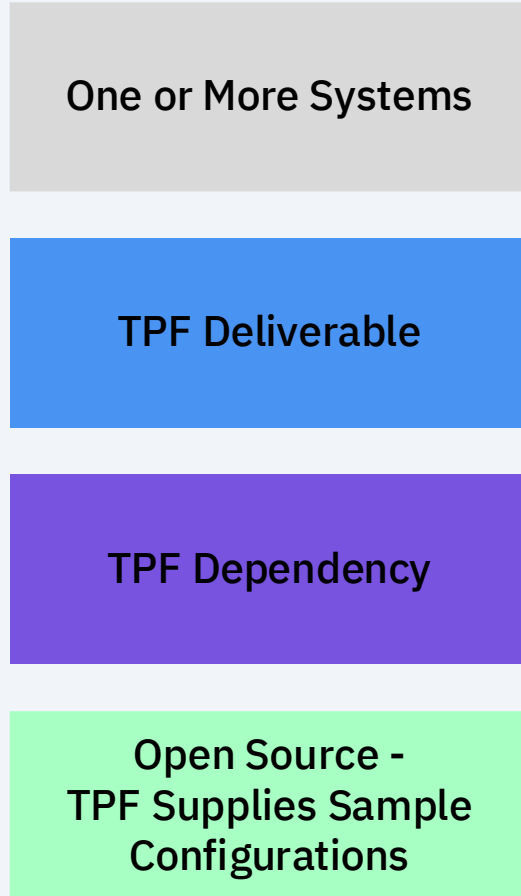
Technical Details

- Real-time Insights Dashboard Starter Kit, customer exploitation:
 - You can implement machine learning and statistical techniques to provide quick insights into system resource usage to speed up problem determination.
 - You can implement alerts and more.
 - If you've implemented the MsgType name-value pair and use the ECB owner name high as the code package, you can leverage the starter kit with your production code without making any changes.

Technical Details

Starter Kit with Replay Script Experimentation Environment

Color Key:



Technical Details

Starter Kit with Real-time Runtime Metrics Collection Experimentation Environment

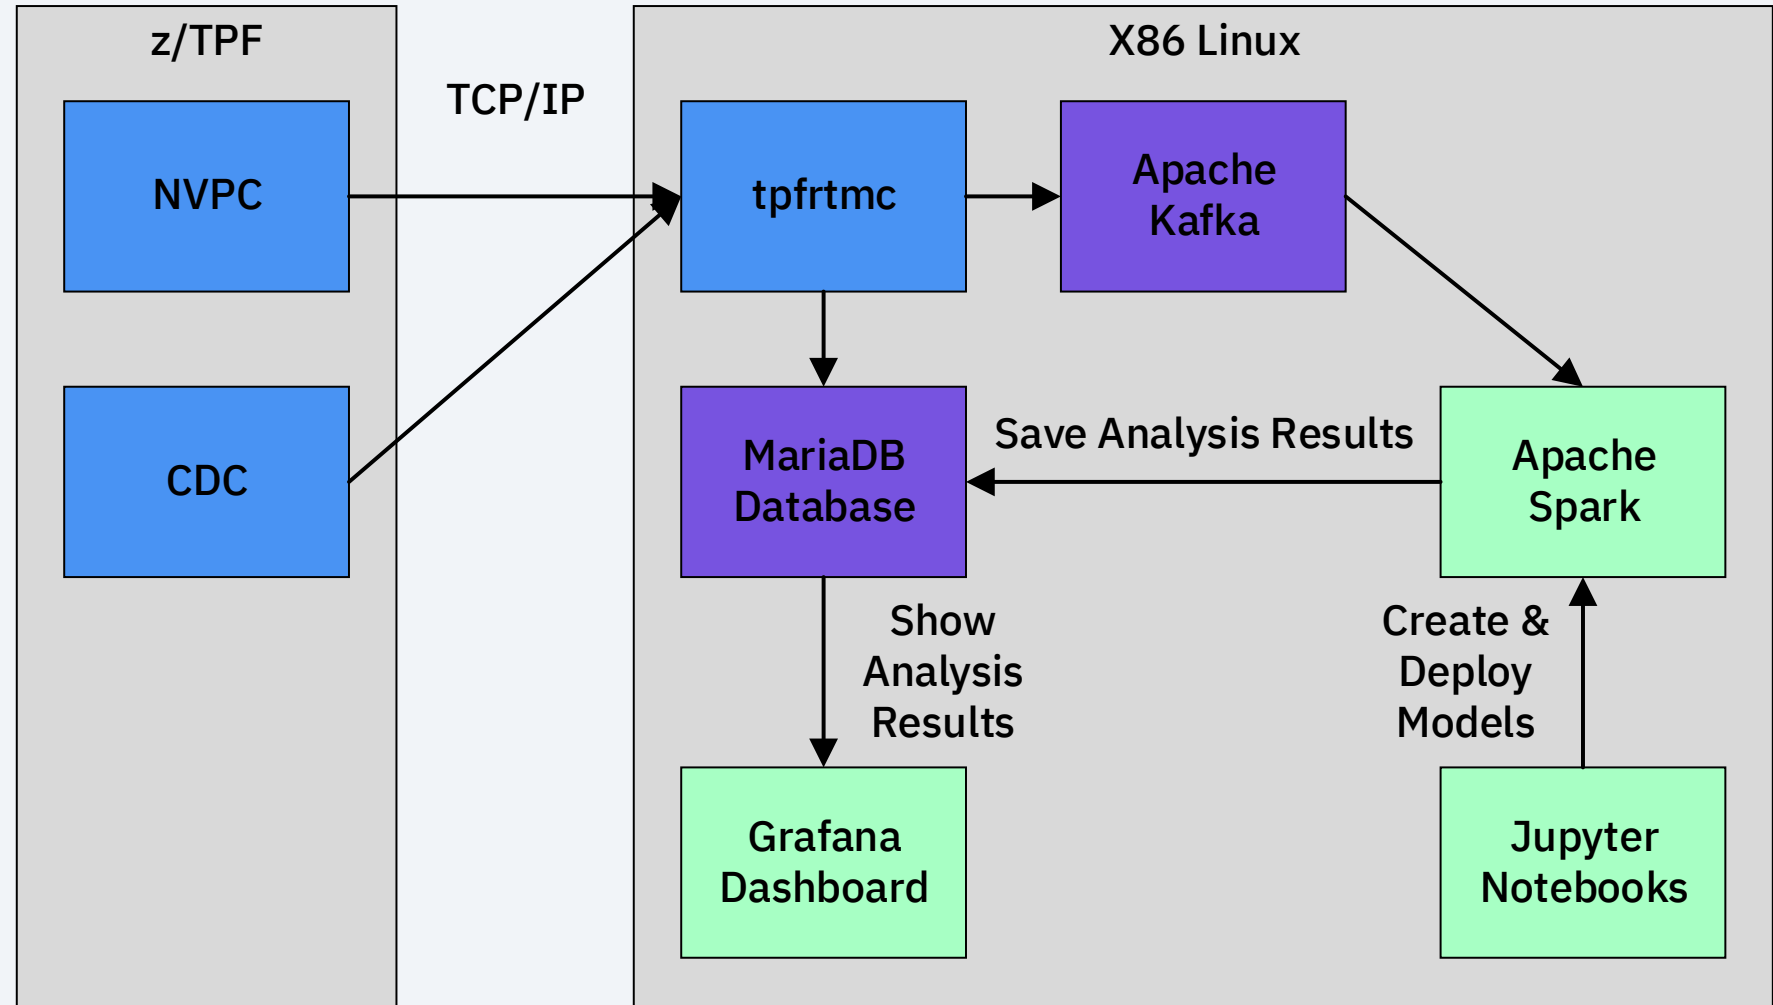
Color Key:

One or More Systems

TPF Deliverable

TPF Dependency

Open Source -
TPF Supplies Sample
Configurations



Technical Details

Starter Kit with Real-time Runtime Metrics Collection Production Environment

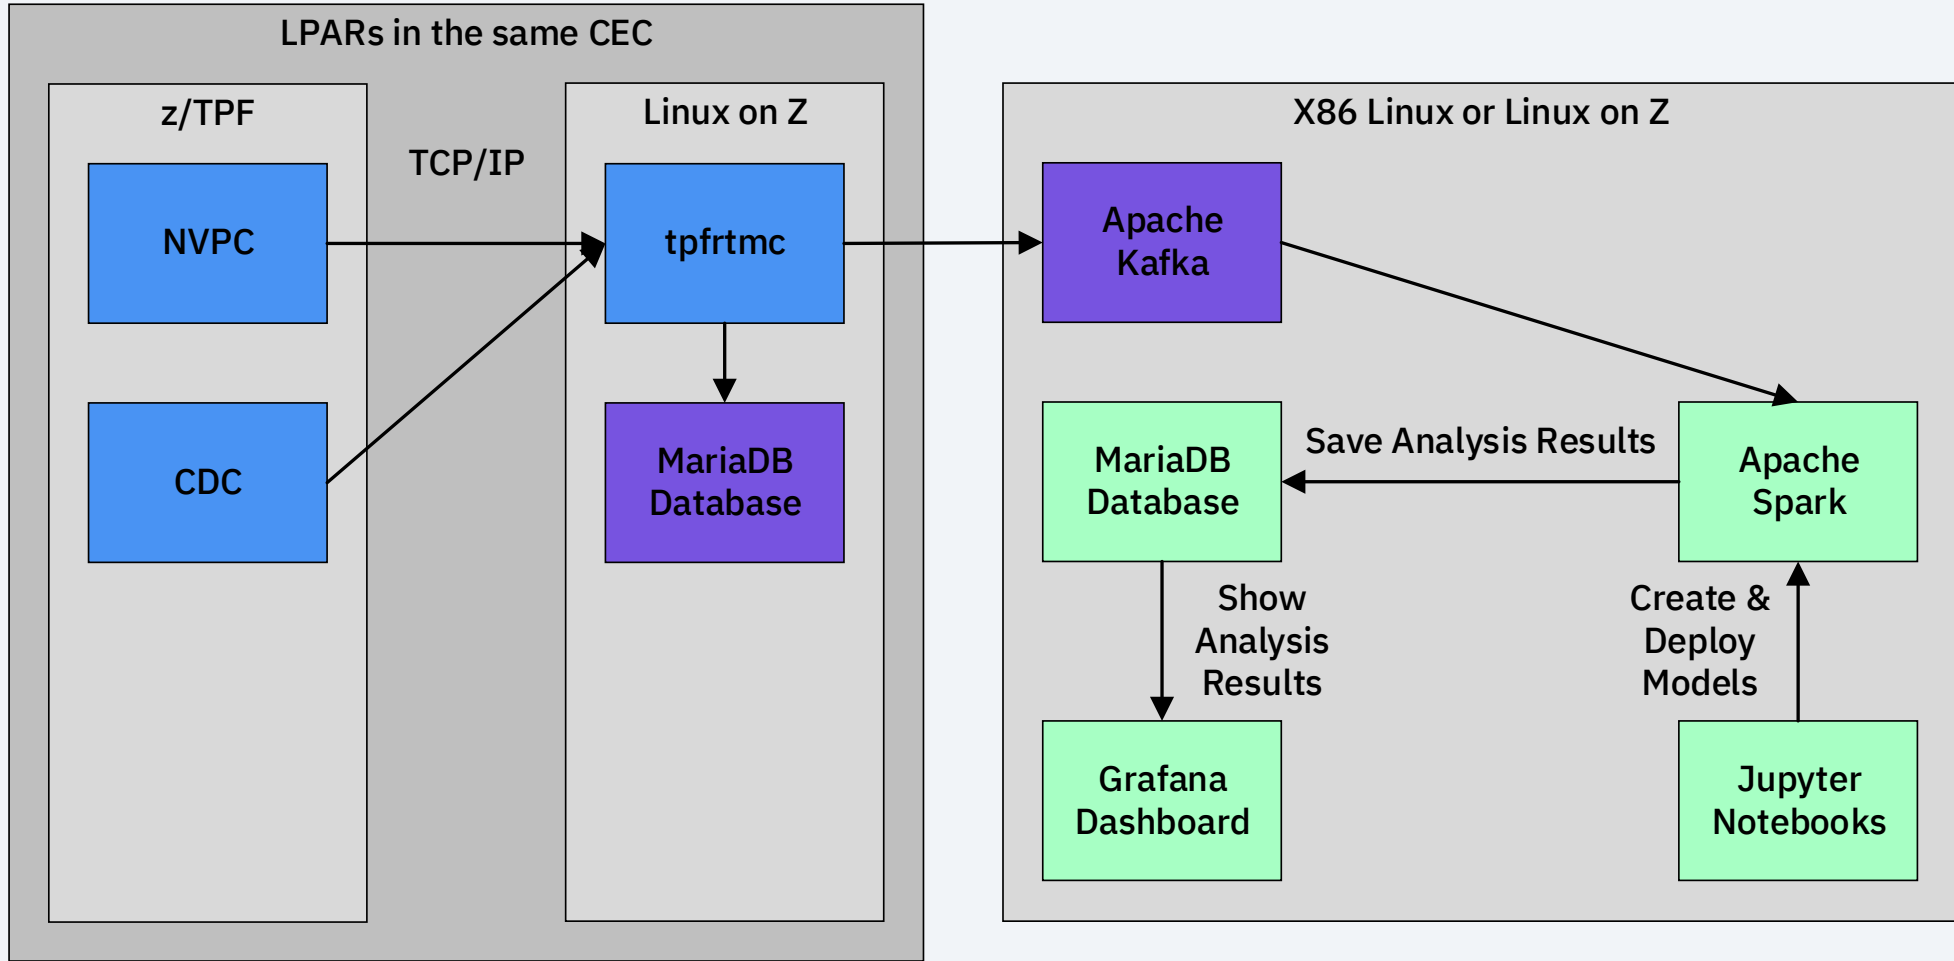
Color Key:

One or More Systems

TPF Deliverable

TPF Dependency

Open Source -
TPF Supplies Sample
Configurations



Installation:

Sets up and configures all components for the Starter Kit with Replay Script Experimentation Environment. Easily extended to use Real-time Runtime Metrics Collection.

1. Install docker and docker-compose
2. `tar -xf tpf_realtime_insights_dashboard.tar`
3. `./tpf_prepare_configurations.sh`
4. `docker-compose up -d`
5. `docker exec -it tpf-mariadb /bin/bash`
`/home/tpftrmc/SQL/tpf_setup_mariadb.sh`
6. `./tpf_create_kafka_topics.sh`

Demo

Running Replay Script:

1. `./tpf_run_spark.sh`
2. `./tpf_start_replay_script.sh scenario_directory`
3. Open Grafana: `http://your.server.name.com:3000`

Conclusion

- Real-time Runtime Metrics Collection
 - PJ45657 is available.
 - 80% of CDC data types are available in the initial release.
 - PJ46185 is now available and adds 2 new types of CDC data.
- Real-time Insights Dashboard Starter Kit
 - Available as of December 2019.
 - Updated with corresponding changes for PJ45657.
 - Download from z/TPF Tools download web page:
<https://www.ibm.com/support/pages/ztpf-real-time-insights-dashboard-starter-kit>

Value Statement

Real-time Runtime Metrics Collection and Real-time Insights Dashboard Starter Kit

Coverage programmers can leverage data science insights into system state and application trends in real time to prevent impacts to service level agreements.

What's next?

- Real-time Runtime Metrics Collection Enhancements
 - Additional data feeds and metrics.
 - Such as those required for Monitoring REST Services.
- Insights Dashboard Starter Kit Enhancements
 - Collect data for machine learning.
 - Implement machine learning models.
 - Implement predictive models.
 - Change point, trend and other forms of detection and alerts.
- Let us know if your interested in being a sponsor user for any upcoming efforts. jwisniew@us.ibm.com danielle.tavella@ibm.com

What's next?

- z/TPF Message Analysis Tool
 - At the TPFUG last year, customers emphasized their need for a performance analyzer replacement.
 - The Message Analysis Tool is a very different tool that does much more than just time analysis.
 - We are actively working on this tool now!
 - We are working with customers to define features and priorities in more depth.

- Let us know if your interested in being a sponsor user for any upcoming efforts. jwisniew@us.ibm.com
danielle.tavella@ibm.com

Virtual TPFUG Q&A

Summary of Q&A from the virtual TPFUG event:

Question	Answer
Q: Is Name-value pair metrics only available through restful services type communication ?	A: Name-value pair metrics are available without restful services. Use the tpf_nameValueLocalSet api to set any name-value pair you like. For REST services, the system sets some name-value pairs automatically for you based upon the input to the REST service as described here: https://www.ibm.com/support/knowledgecenter/en/SSB23S_1.1.0.2020/gtpa2/sysnamevaluepair.html
Q: does the analytics pipeline have to be installed on x86 Linux? Can it be Linux on Z?	A: The analytics pipeline including the offline tpftrmc utility can be installed on Linux on Z.
Q: So to see Dashboard, a user needs access to Grafana?	A: In our sample implementation, a user utilizes Grafana Dashboards through a web interface to access the data processed by tpftrmc. However this is just a sample implementation. Data could be read from Apache Kafka, analyzed and managed however you like.

Virtual TPFUG Q&A

Summary of Q&A from the virtual TPFUG event:

Question	Answer
<p>Q: For production environments, is it a hard requirement to have Linux on Z and z/TPF LPAR as same CEC? Can we push data from different LPARs to Linux on Z on different CEC.</p>	<p>A: No, there is no hard requirement to have them on the same CEC. But there are benefits if you do so such as less TCP/IP overhead. With the high speed connector, real-time runtime metrics collection is highly configurable including fall back instances. Currently, the tpftrmc offline utility can only receive data from a single LPAR but we expect this restriction to be removed at some point in time.</p>
<p>Q: What is the TPF dependency for MariaDB Database?</p>	<p>A: There is no dependency on the MariaDB database for the online z/TPF system. The offline utility, tpftrmc, depends on MariaDB to do the analysis. But for real-time runtime metrics collection, the data is transient and is not preserved.</p>
<p>Q1: Will tpftrmc run on x86 linux? Q2: So, NVPC can send info via HSC to an "external" instance of tpftrmc?</p>	<p>A1: Yes, tpftrmc will run on x86 or Linux on Z A2: Yes, Real-time runtime metrics collection sends data via the HSC to an "external" instance of tpftrmc offline. Note that name-value pair collection (ZCNVP) does not use the HSC, it puts the data on tape which is then post-processed. Real-time runtime metrics collections (ZRTMC) uses the HSC to send the data.</p>


Virtual TPFUG Q&A

Summary of Q&A from the virtual TPFUG event:

Question	Answer
<p>Q1: On the message analysis dashboard, is there any correlation between the orange squares in the widget and the Lines in the Chart? ie Coef-Rt is related to CPU Line or something line that...</p>	<p>A1: The correlation results are related to the data being shown in the chart. However, as described, some of the analysis results (The orange squares) will be improved upon in an upcoming release. Yes, Coef-Rt is related to the CPU line. To further elaborate, each correlation coefficient in each panel has a different meaning. For example, one panel correlates the message rates for each message type to system CPU utilization. If the correlation breaks a threshold, the coef cell is changed to orange. See each panel in the starter kit and education dashboards for more information.</p>

Thank You

Questions? Comments?





Trademarks

IBM, the IBM logo, ibm.com and Rational are trademarks or registered 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.

Notes

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-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.