Name Value Pair Collection (NVPC)

Eric Tryon
TPF Developer

IBM z/TPF

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

1. What is Name Value Pair Collection?
    - Purpose and Value
    - Current State vs Future State
    - NVPC in the Big Picture
    - NVPC Results
    - Use Cases
    - Horizontal NVPs vs Vertical NVPs
2. NVPC Architecture
    - NVP Collection Command
    - TPFRTMC Server
    - ADI User Interface
3. Moving Forward
    - Planned Release Schedule
    - Prepare for NVPC!

What is NVPC?

- Name value pair collection identifies:
  - The messages that were processed by the system.
  - The average resources used to process the different messages.
  - The code packages or phases of processing used to process messages.
  - The average resources used by the code packages or phases of processing.

Current State

Carol the coverage programmer receives a call from operations regarding a degradation in system performance.

By way of time consuming investigation, clever techniques and raw determination, she finds a set of modules is performing particularly poorly for a given customer.

Carol does not have the tools to quickly identify unanticipated sources of resource consumption.

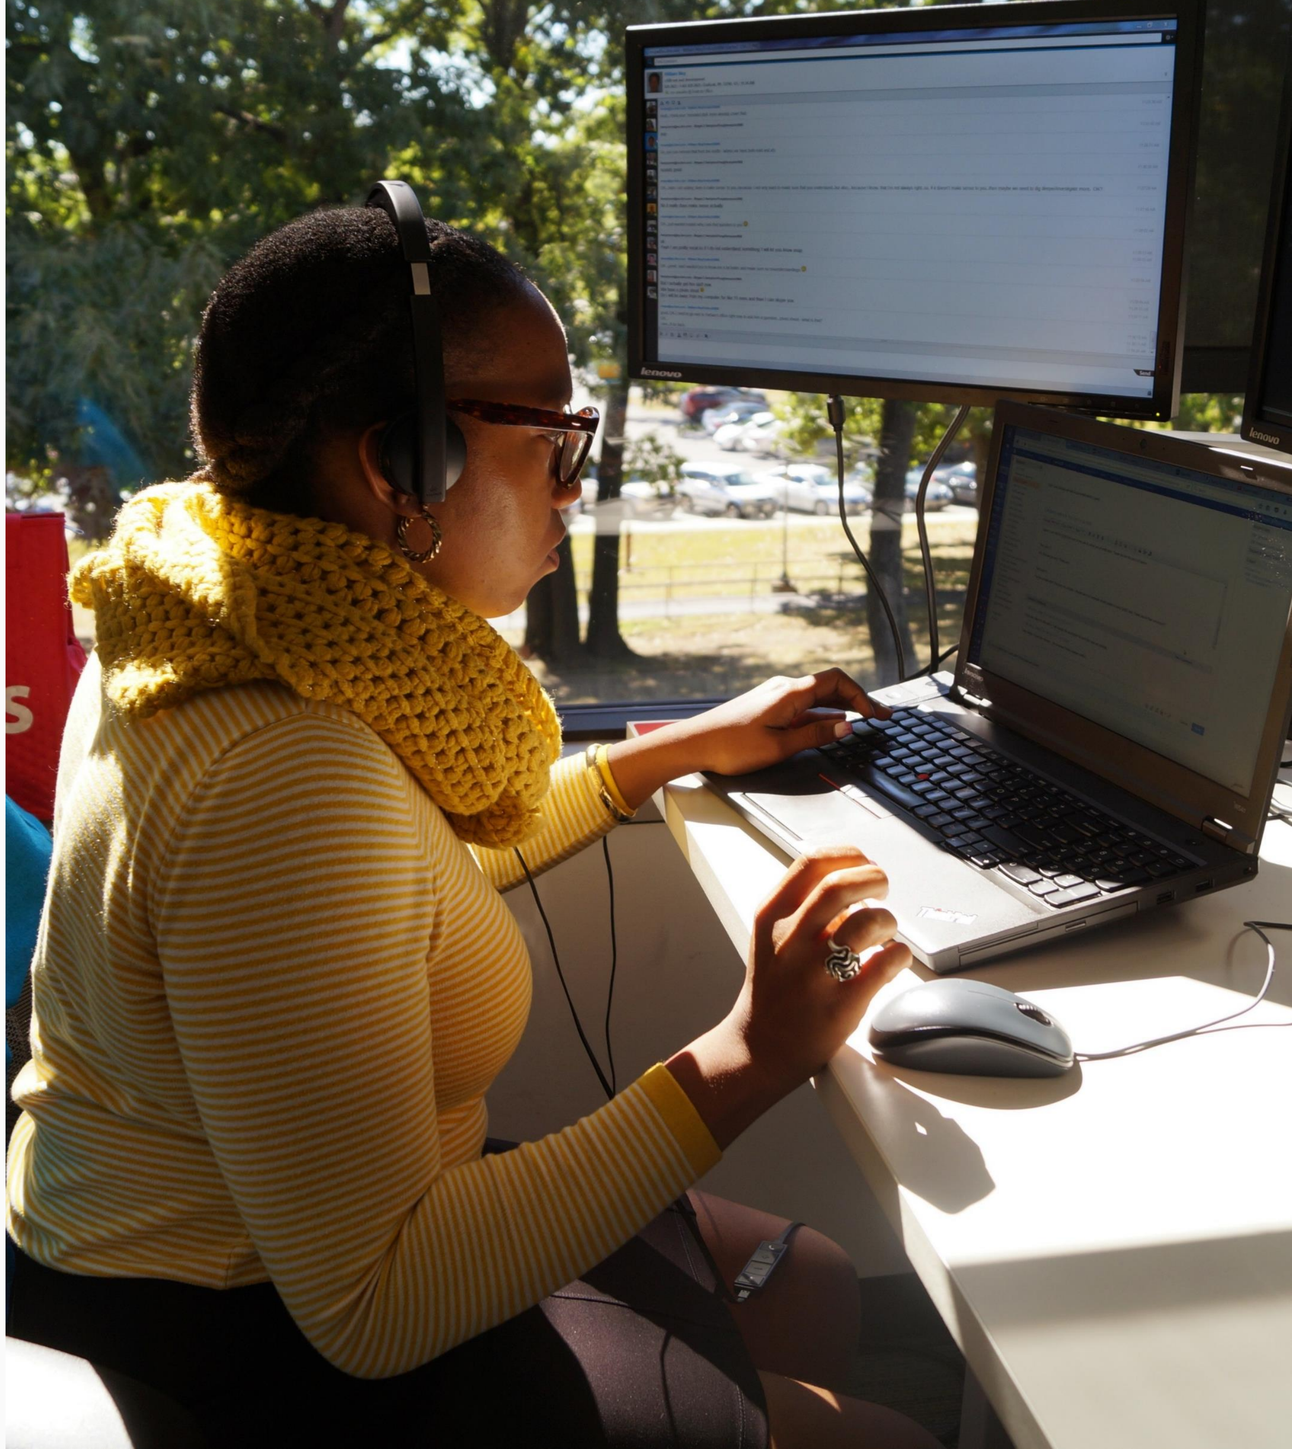

Future State

Carol the coverage programmer receives a call from operations regarding a degradation in system performance.

She uses name-value pair collection to quickly identify mobile traffic from a particular customer is doing more IO than normal.

Carol can quickly identify sources of resource consumption enabling her to focus on the causes.

Value Statements

- A coverage programmer can use name-value pair collection to gain new insights into system resource usage and identify the source of problems in as little as 1/20 of the time previously required.
- A capacity planner can use name-value pair collection to determine the additional physical assets required for the expected message rate growth of a specific service given new resource usage metrics.
- An application architect can use name-value pair collection to identify inefficient code packages that can be refactored to improve system performance.

NVPC in the Big Picture

- Use Data Collection/Reduction, Continuous Data Collection (CDC) and other tools to understand system resource usage across the system.
- Use Owner Name Collection, Software Profiler (ZTRAP) and other tools to understand resource usage in broad strokes or for particular code areas.
- Use Name Value Pair Collection to “zoom in” and understand resource usage at both the
  - The message level.
  - The code packages or phases that implement the message processing.

NVPC Results in ADI

- Application Delivery Intelligence (ADI) is z DevOps tooling which provides analysis, insights and trends into application structure, test coverage and run time metrics.
- For NVP collection, ADI will provide:
  - Ability to view and compare the collection results.
  - Analysis of collections.
  - Long term trends.
  - And more.

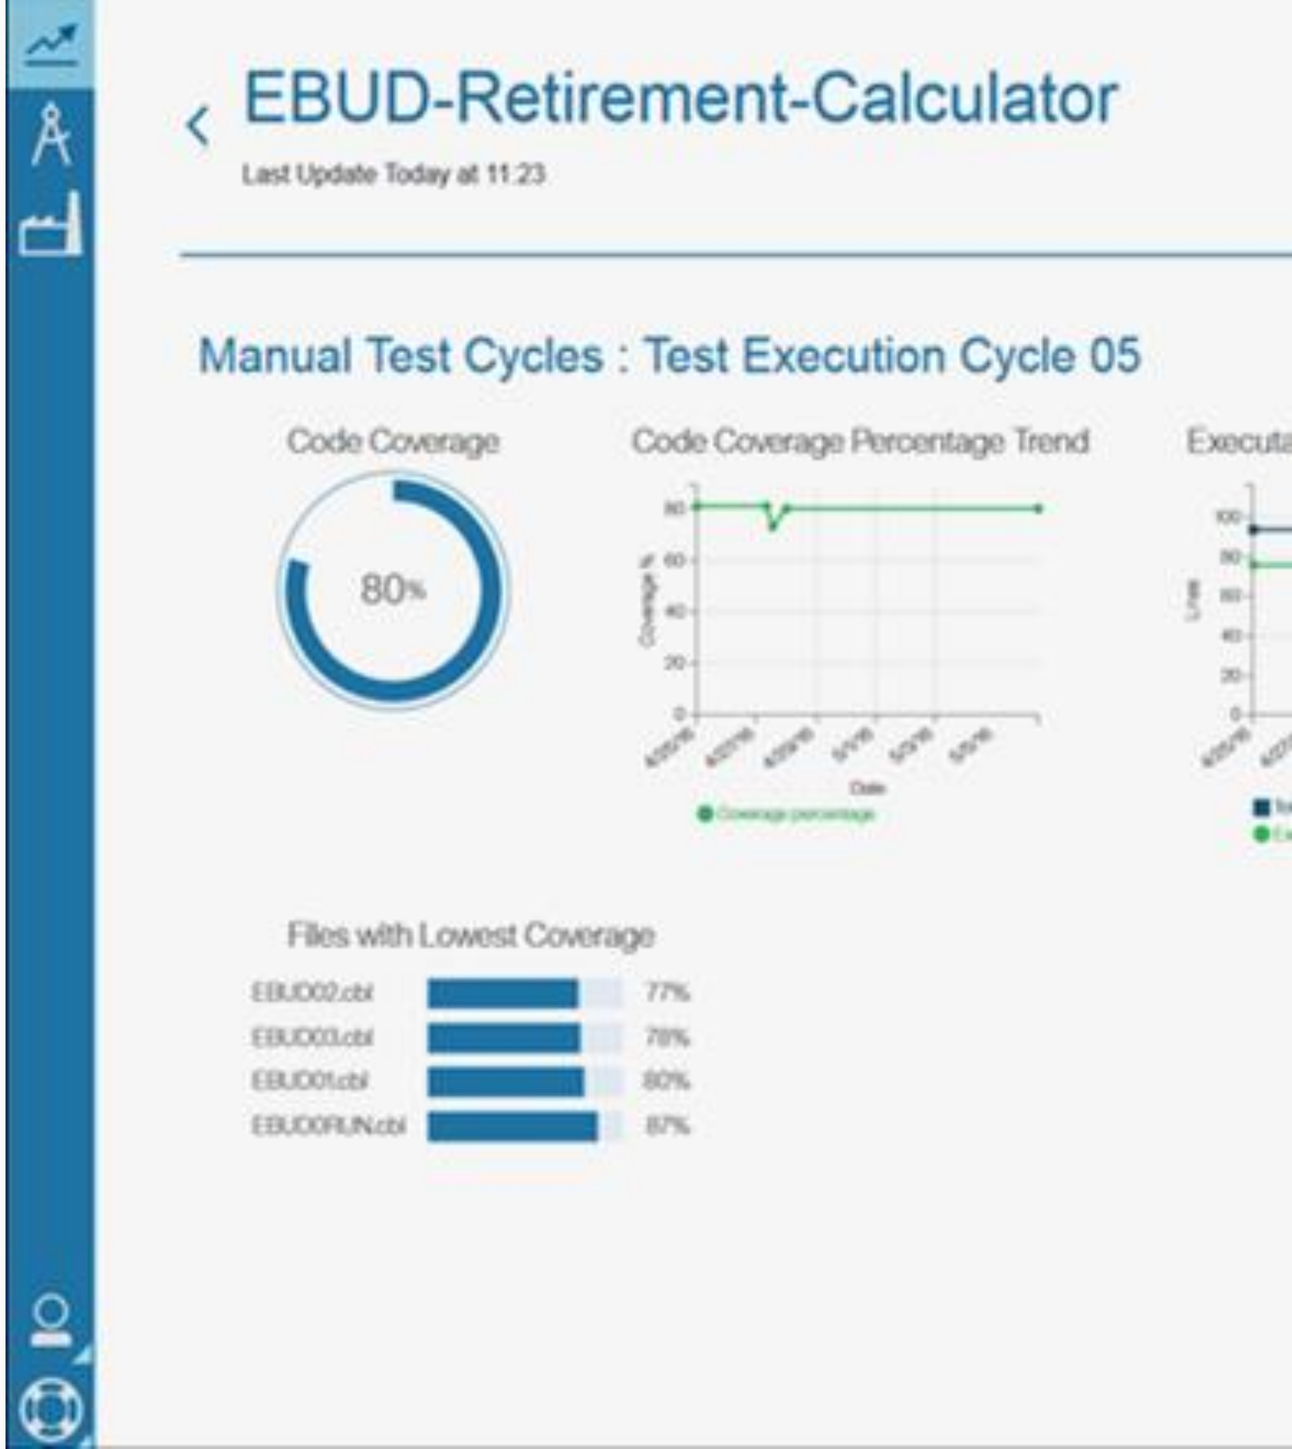

NVPC Use Cases

Daily Usage

- Run NVP collection on TPF daily around the same time you run data collection (ie at peak).
- View the NVP collection results in ADI.
- Operations can quickly interpret graphs that indicate:
  - Message mix compared to one or more previous collections.
  - Resources used per message compared to one or more previous collections.
- Operations, coverage or etc. can dig deeper into analysis, trends and a table of data.

NVPC Use Cases

System Performance Issue Usage

- A coverage programmer could leverage an existing NVP collection for their investigation.
- This collection is viewed in ADI.
- ADI allows you to leverage different relationships between the NVPs to reveal different insights into the collected data. For example:
    - By MsgType and Source (ie mobile, terminal, etc)
    - Filter on MsgType=Booking and by OwnerNameHi (code package) and Source(ie mobile, terminal, etc)
- Analysis, trends and the data provides a variety of insights.

NVPC Use Cases

Capacity Planner Usage

- A capacity planner could leverage an existing NVP collection for their investigation.
- This collection is viewed in ADI.
- ADI allows you to leverage different relationships between the NVPs to reveal different insights into the collected data. For example:
    - By Source (ie mobile, terminal, etc) and MsgType
- Analysis, trends and the data provides a variety of insights.
- Understanding the resources currently used by a Source, MsgType or etc in light of the current trends and future projections allows for more accurate hardware estimates.

Horizontal vs Vertical

- Think of name value pairs as describing the entire message processing from start to finish. The horizontal. The whole.

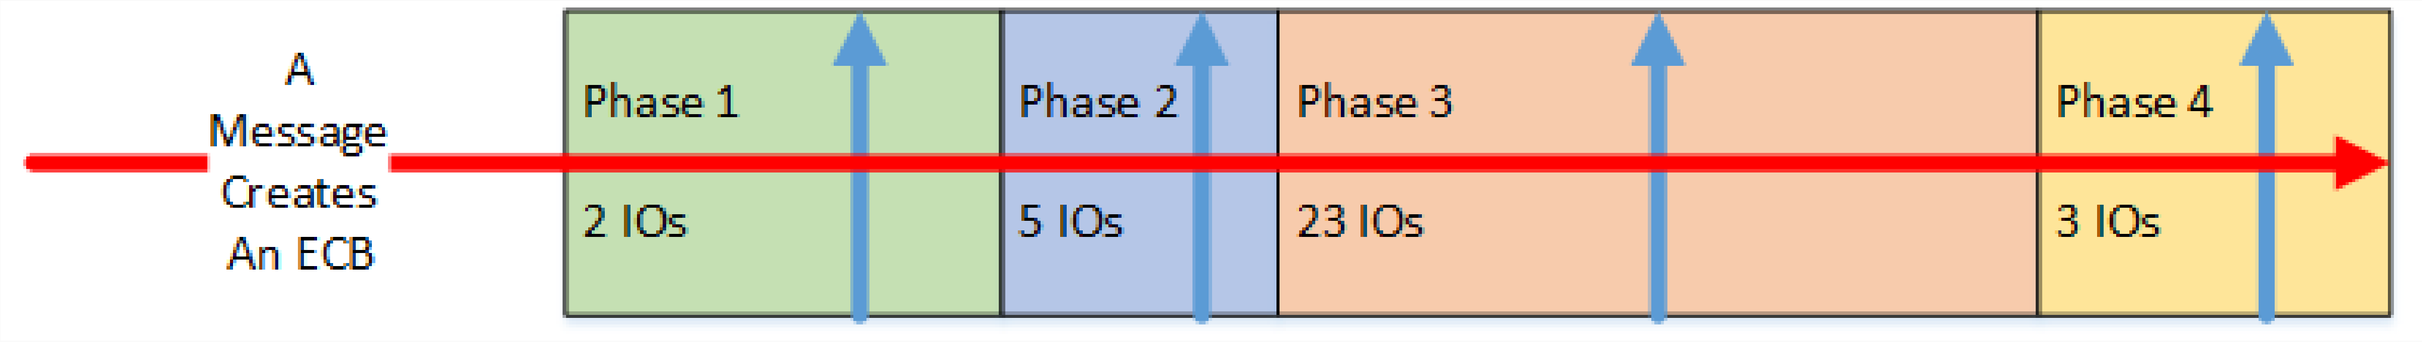

Horizontal vs Vertical

- Name value pairs describe messages in the system in multiple categories as you require. For example:
  - type of the message, why is it here
  - the source of the message, how did it get here
  - the purpose of the message, even describe input values
- You set name value pairs at the beginning of processing a message and describe the message as needed.
- NVPs generally should not be changed after setting.
- NVPs are inherited by child ECBs and so resource usage analysis is for the entire unit of work.

Horizontal vs Vertical

- Think of ECB owner names as describing the code packages or phases of processing throughout the processing of the message. The vertical. The piece parts that make up the whole.

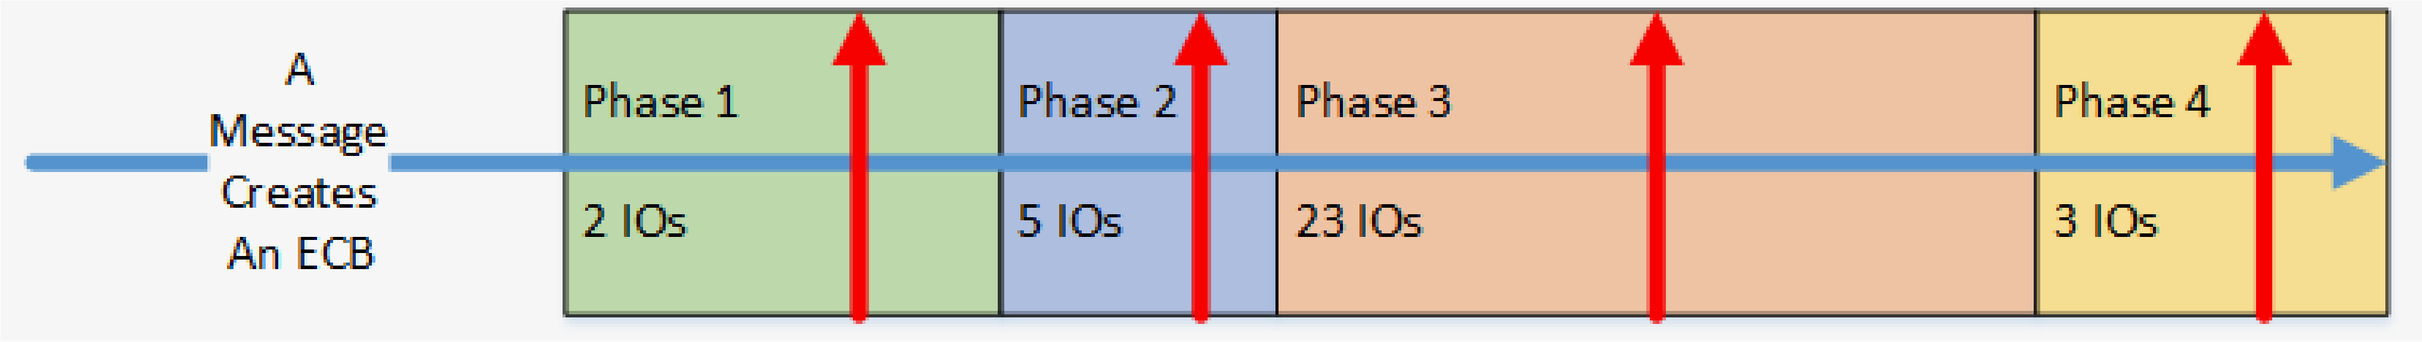

Horizontal vs Vertical

- ECB Owner names describe the pieces that make up the message. For example:
  - Code package or
  - Phases of processing
- Set at the beginning of message processing.
- Change at the primary vertical boundaries such as entering a new phase, entering a new code package, etc.
- Do NOT change the owner name too frequently as the overhead of collection may become noticeable in your system resource usage. For example, do NOT set ECB owner name at every entry point, function call, etc.

More Details

2016 TPFUG operations session: https://www-01.ibm.com/software/htp/tpf/tpfug/tgf16/TPFUG_2016_OPER_ADI.pdf

2016 TPFUG education session: https://www-01.ibm.com/software/htp/tpf/tpfug/tgf16/TPFUG_2016_EDUC_ADI.pdf

Architecture

- NVP Collection occurs on TPF.
- Results are written to tape for minimal system impact during collection.
- TPFRTMC server processes the tape file and takes care of TPF specific details.
- ADI Server analyzes the collection results to create trends, insights, and tables of data.
- User investigates ADI results through a web browser interface.

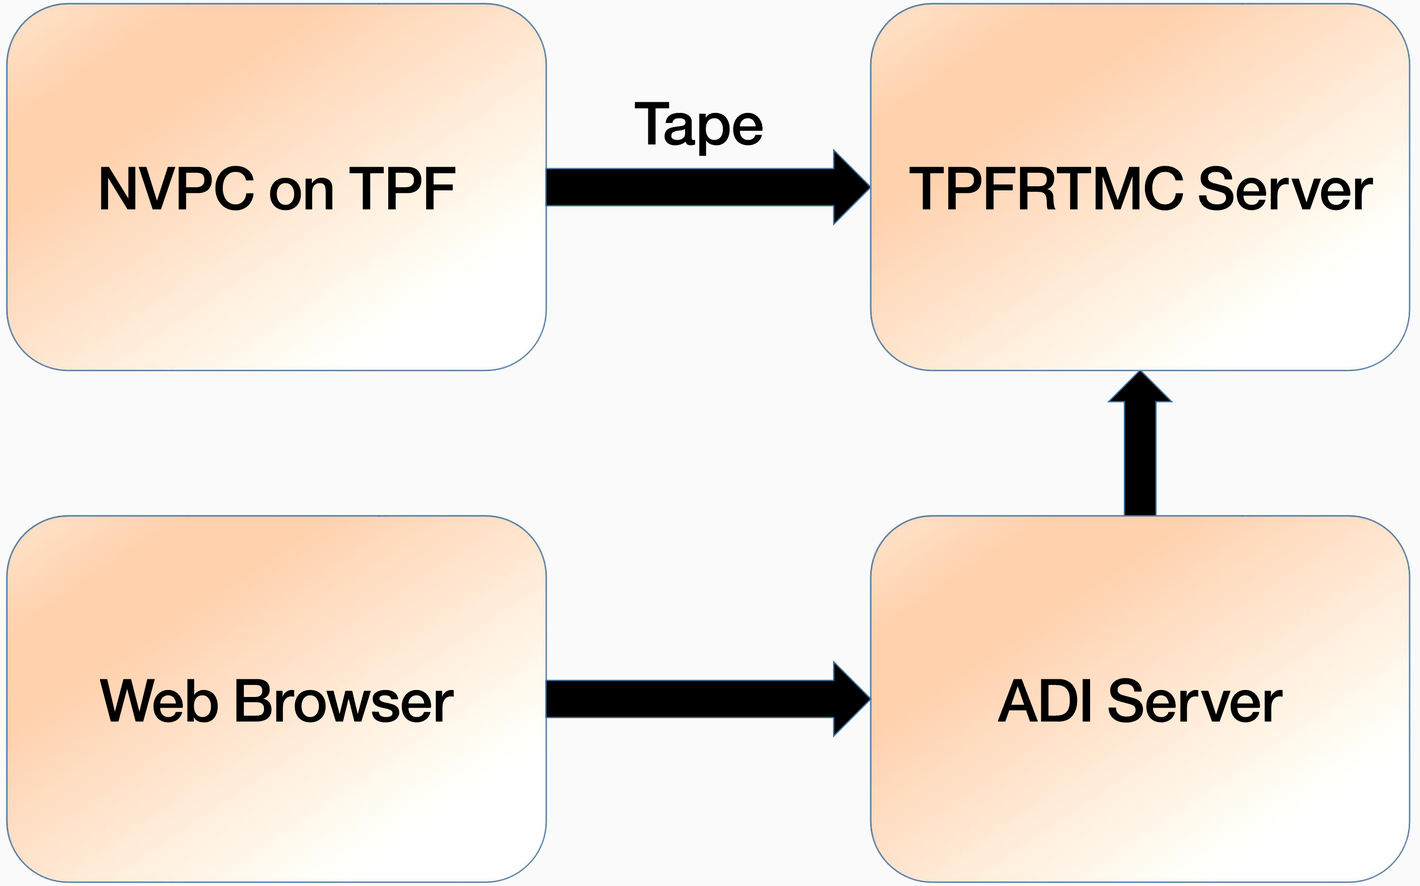

Architecture

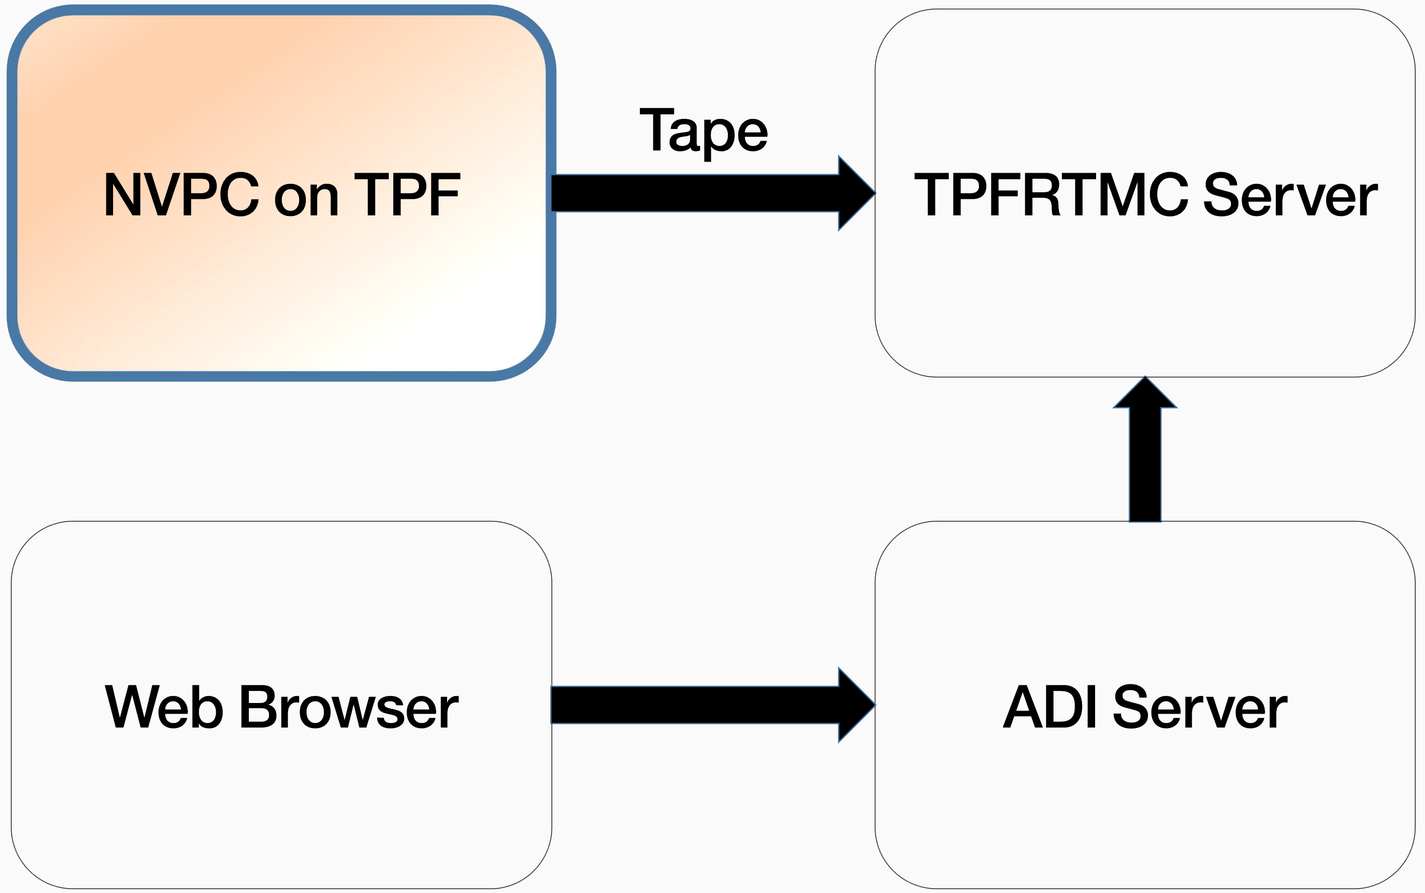

NVP Collection Command

1. Mount a blocked general tape on the system for name-value pair collection.
2. Create a name-value pair policy file on the TPF file system such as: /tmp/my.cfg. The policy tells TPF which name-value pairs to write out with the metrics data.

```
POLICY_VERSION=1
MsgType,horizontal
MsgSource,horizontal
#END OF THE PROFILE  FILE
```

NVP Collection Command

3. Start the name-value collection with the new TPF command:

```
ZCNVP START          SAMPLE
                     MSGSKIP-1000
                     OWNERSKIP-5
                     TAPE-za1
                     POLICY-/tmp/myCol.txt
                     HOWLONG-10
```

Collects metrics at owner name change and ECB exit.
Skips 1000 messages before collecting an entry (range is 0 to 99999).
Skips 5 owner name change events before collecting an entry (range is 0 to 9).
Collects for 10 minutes (range is 1 to 999 minutes).

NVP Collection Command

A variation is available

ZCNVP START

EXIT
MSGSKIP-1000
TAPE-za1
POLICY-/tmp/myCol.txt
HOWLONG-10

Collects metrics at only at ECB exit.
Skips 1000 messages before collecting an entry (range is 0 to 99999).
Collects for 10 minutes (range is 1 to 999 minutes).

NVP Collection Command

4. Collection completes or you can use the following commands to manage the name-value pair collection that is in progress:

ZCNVP STATUS – indicate if name-value pair is running and how many entries were not captured due to the tape buffers being full.

ZCNVP STOP – Stops the name-value pair collection.

ZCNVP CANCEL – Stops the name-value pair collection and releases the system heap but does not close or dismount the tape. The data on the tape will not be able to be post processed.

Architecture

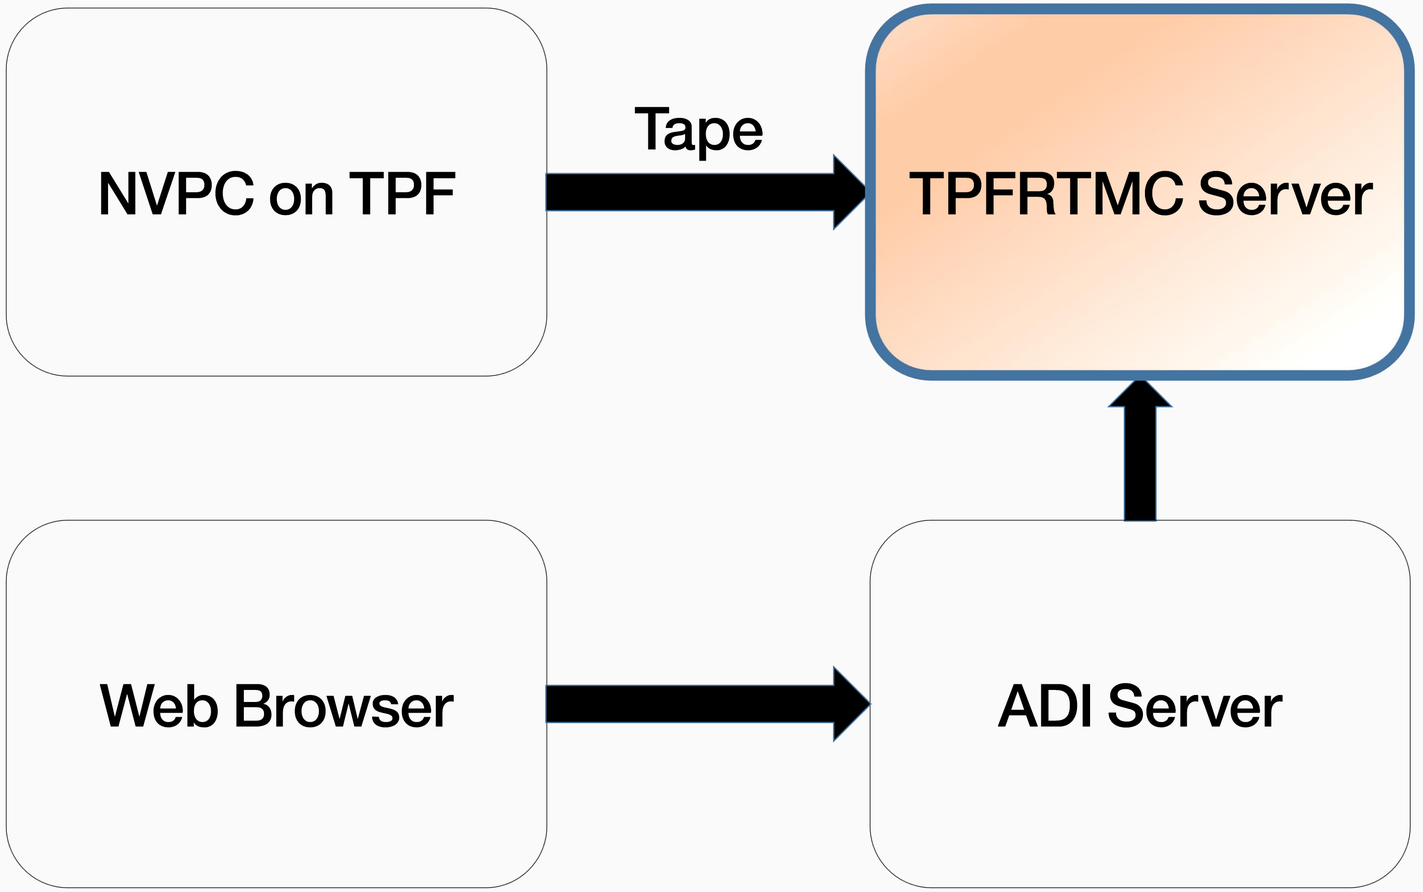

TPFRTMC Server

- TPFRTMC reads the contents of the NVP collection tape and hides the TPF system specifics from the ADI Server.
- The tape produced by online NVP collection must be converted to a binary file. This can be accomplished by FTPing the tape contents or leveraging the pptapetofile utility.
- The binary file is moved to a directory where the TPFRTMC server is configured to monitor.
- Besides the initial administration of the TPFRTMC server and transfer of the binary file, users do not interact any further with the TPFRTMC server.

Architecture

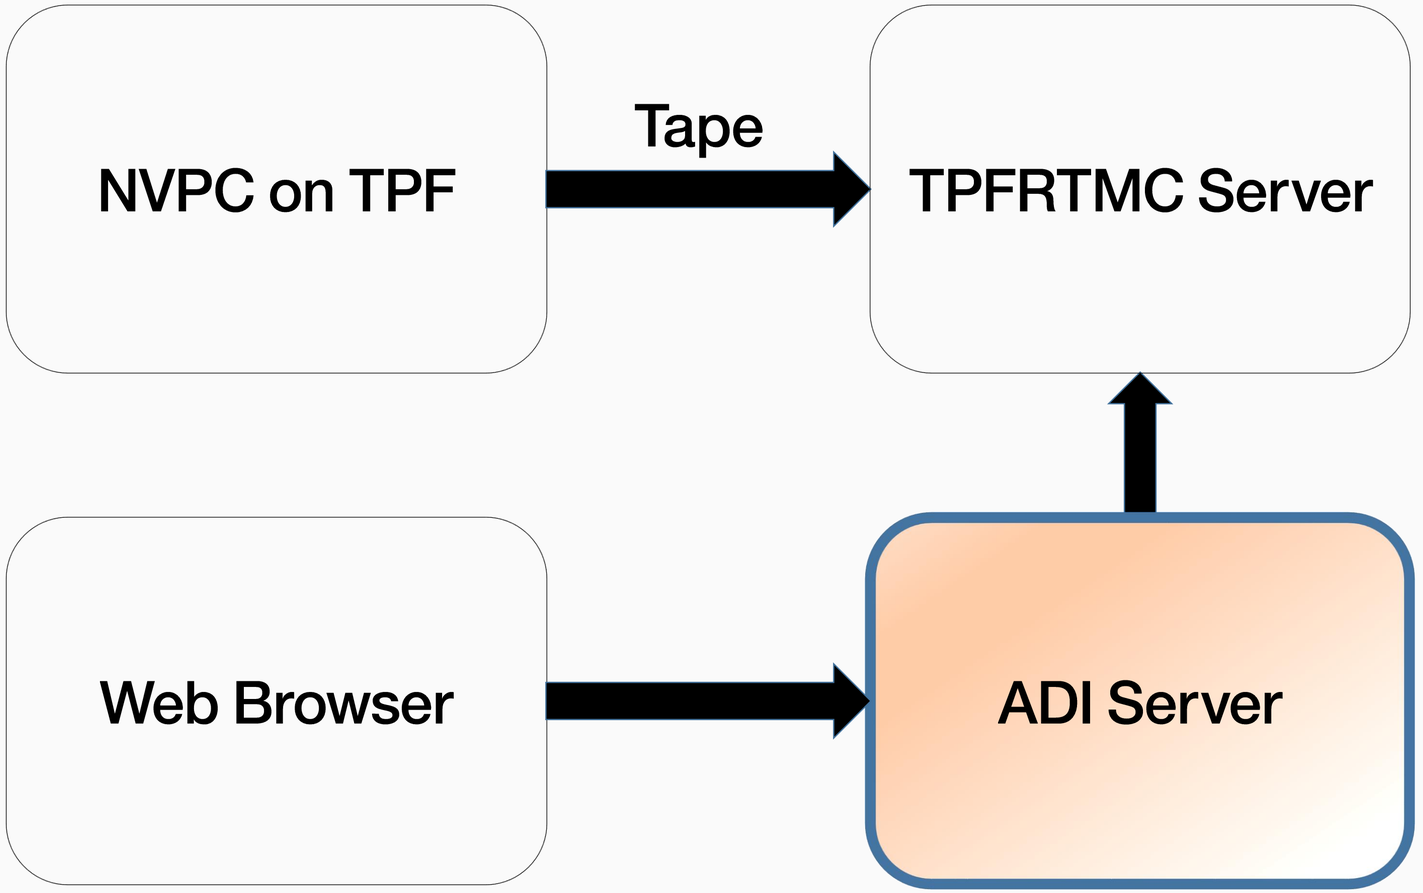

ADI User Interface

- We’d love to show and discuss the design of the UI with you in our customer playbacks and user testing sessions.
- If you are interested in participating, please let us know.

ADI User Interface

- Teaser UI details
  - Ability to define different renderings of data for scheduled or ad hoc (point in time) collections.
  - “Health check” graphs provide a sense/magnitude of changes at a glance.
  - Outlier Analysis provides a weighted analysis of deltas hidden in the data.
  - Long term trend graphs.
  - Table of data with the ability to compare against past collections.

Planned Schedule

- May 2017 TPF online APAR PJ44321 will be released such that you can:
  - Apply the APAR (CP changes).
  - Create a strategy for NVPs and ECB Owner Names. See 2016 TPFUG education session for more details.
  - Modify your application with NVPs and ECB Owner Names.
- June 2017 Plan to start end to end betas (TPF online through ADI).
- 3Q2017 TPFRTMC Server and ADI GA (All components available).
- 4Q2017 Second deliverable to allow analysis of vertical name value pairs such as ECB Owner Names, prune database, etc.

Action Required: Prepare for NVPC

- Create a strategy for NVPs and ECB Owner Names. See 2016 TPFUG education session for more details.
- Modify your application with NVPs and ECB Owner Names.

THANK YOU

Questions or comments?

Eric Tryon
TPF Developer

IBM z/TPF

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.