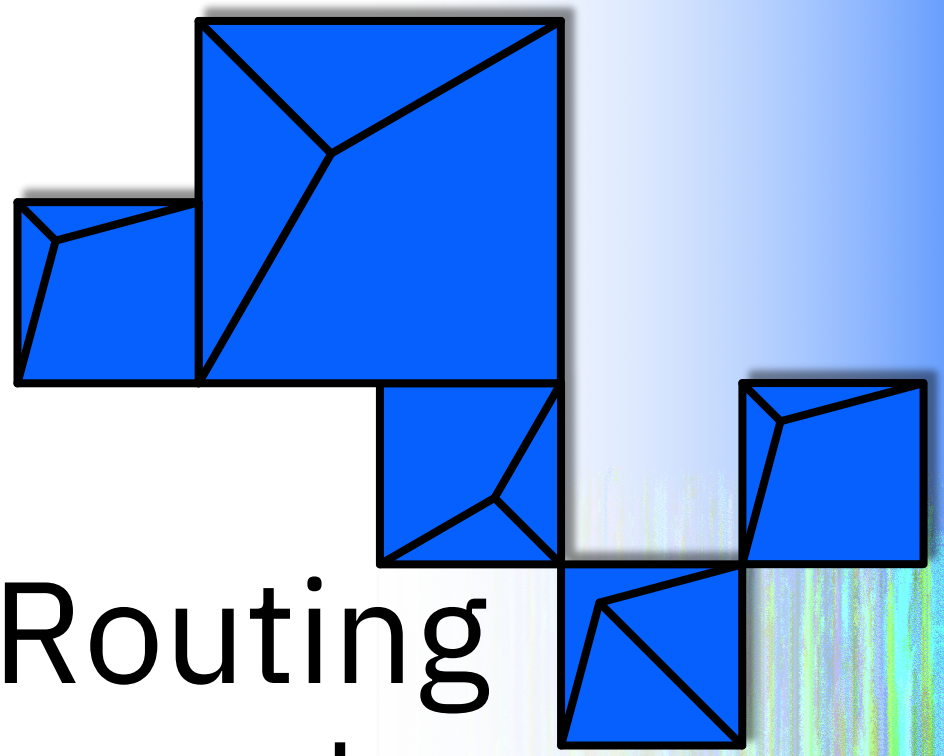

z/TPF MQ Dynamic Routing and Future Enhancements

Jamie Farmer

Agenda

- Background
- Problem Statement
- Pain Points
- Solution
- Technical Details
- Value Statement
- What’s next?

Background

- One of the primary interconnectivity solutions for TPF is using z/TPF MQ.
  - Provides large pipes to transfer data to and from z/TPF.
  - z/TPF MQ can guarantee delivery of messages even in the event of a failure.

- Pushing critical data off of z/TPF through things like z/TPF data events continues to grow.
  - These outbound z/TPF MQ pipes are a critical interface
    - Require high availability
    - Minimize latency to ensure timely delivery of the messages across this interface.

Background

- In z/TPF, when sending messages to a remote queue manager
  - A z/TPF MQ sender channel is created, creating a socket connection to the remote queue manager.
  - A z/TPF MQ transmission queue is created and assigned to the z/TPF MQ sender channel.
  - A z/TPF MQ remote queue definition is created and assigned to the z/TPF Transmission Queue.

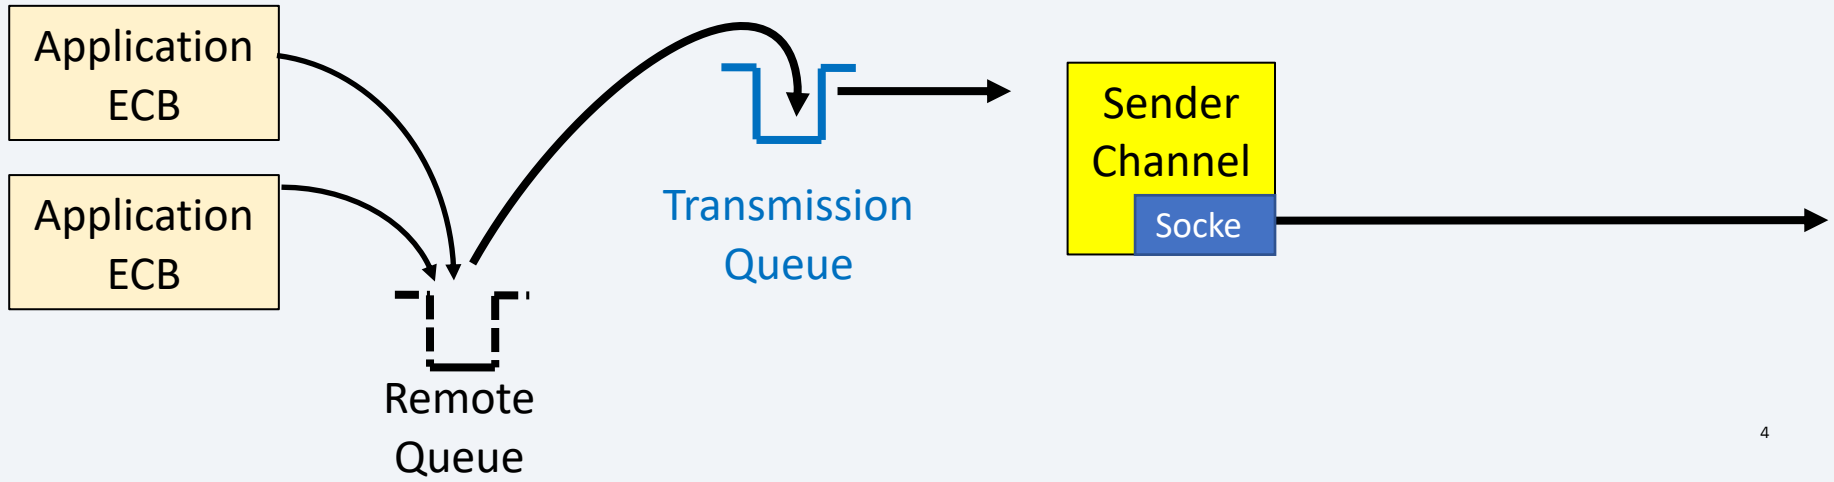

z/TPF MQ Sender Channels

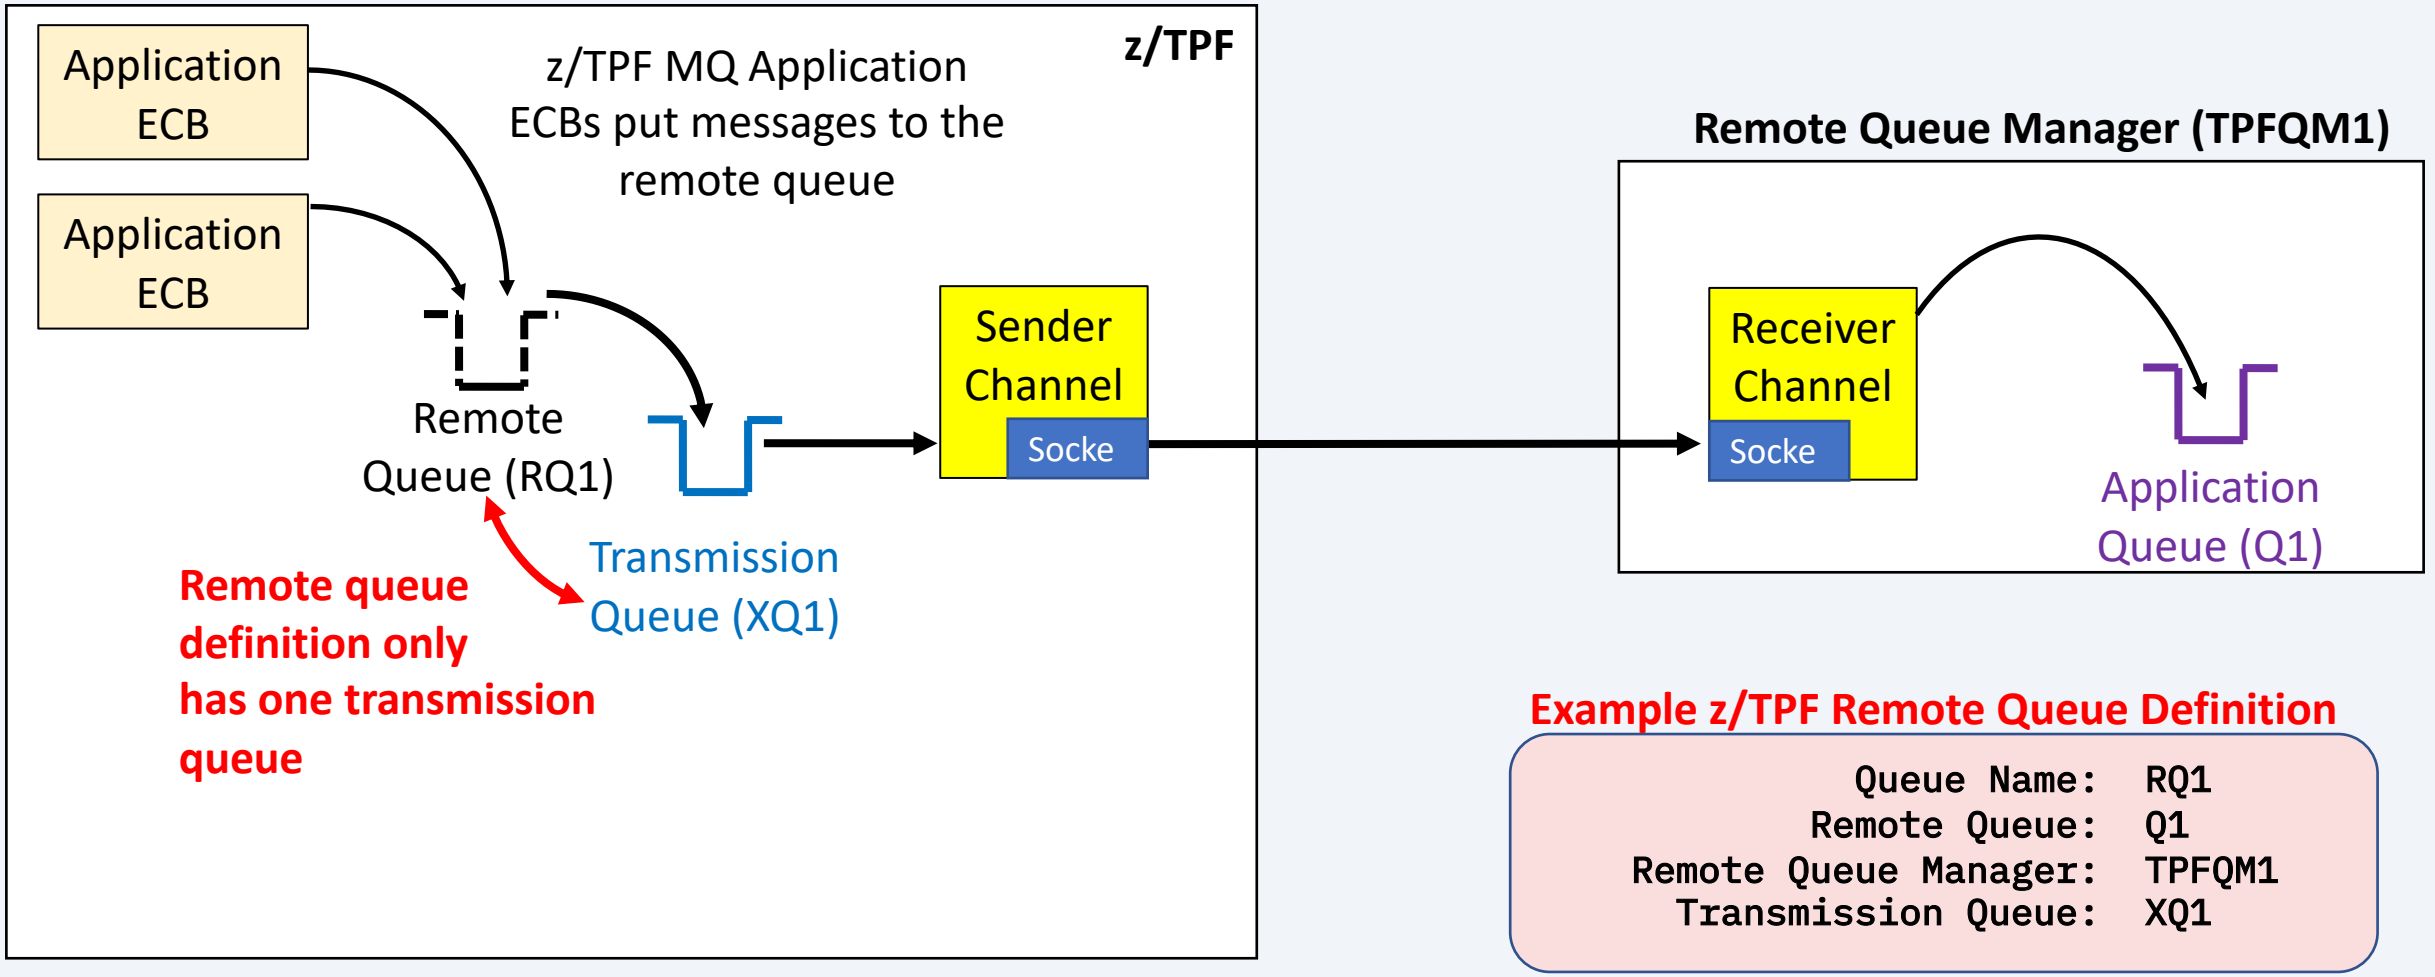

Example z/TPF Remote Queue Definition

```
          Queue Name:  RQ1
        Remote Queue:  Q1
Remote Queue Manager:  TPFQM1
  Transmission Queue:  XQ1
```

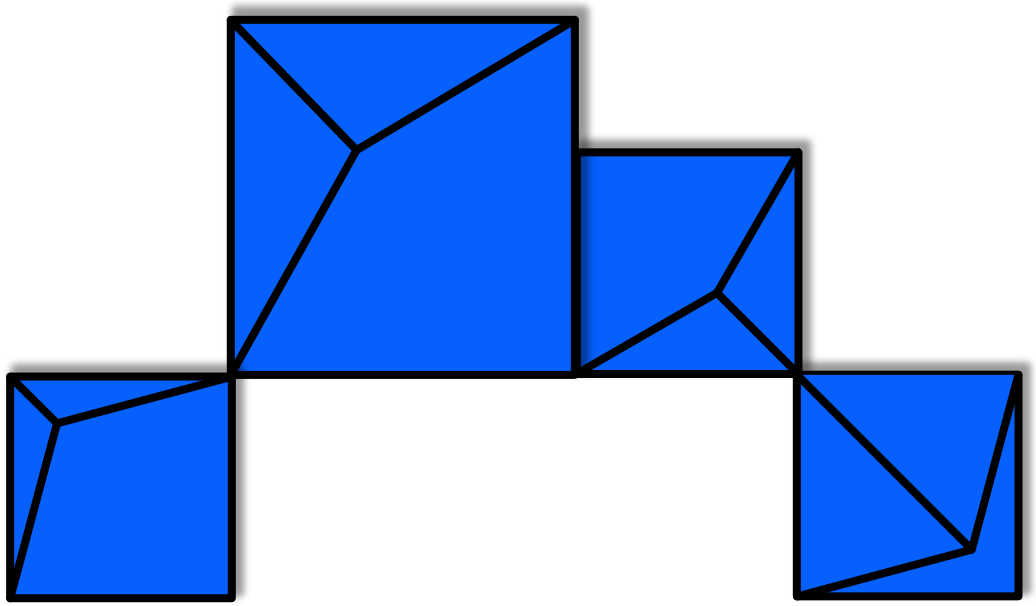

Problem Statement

Problem Statement

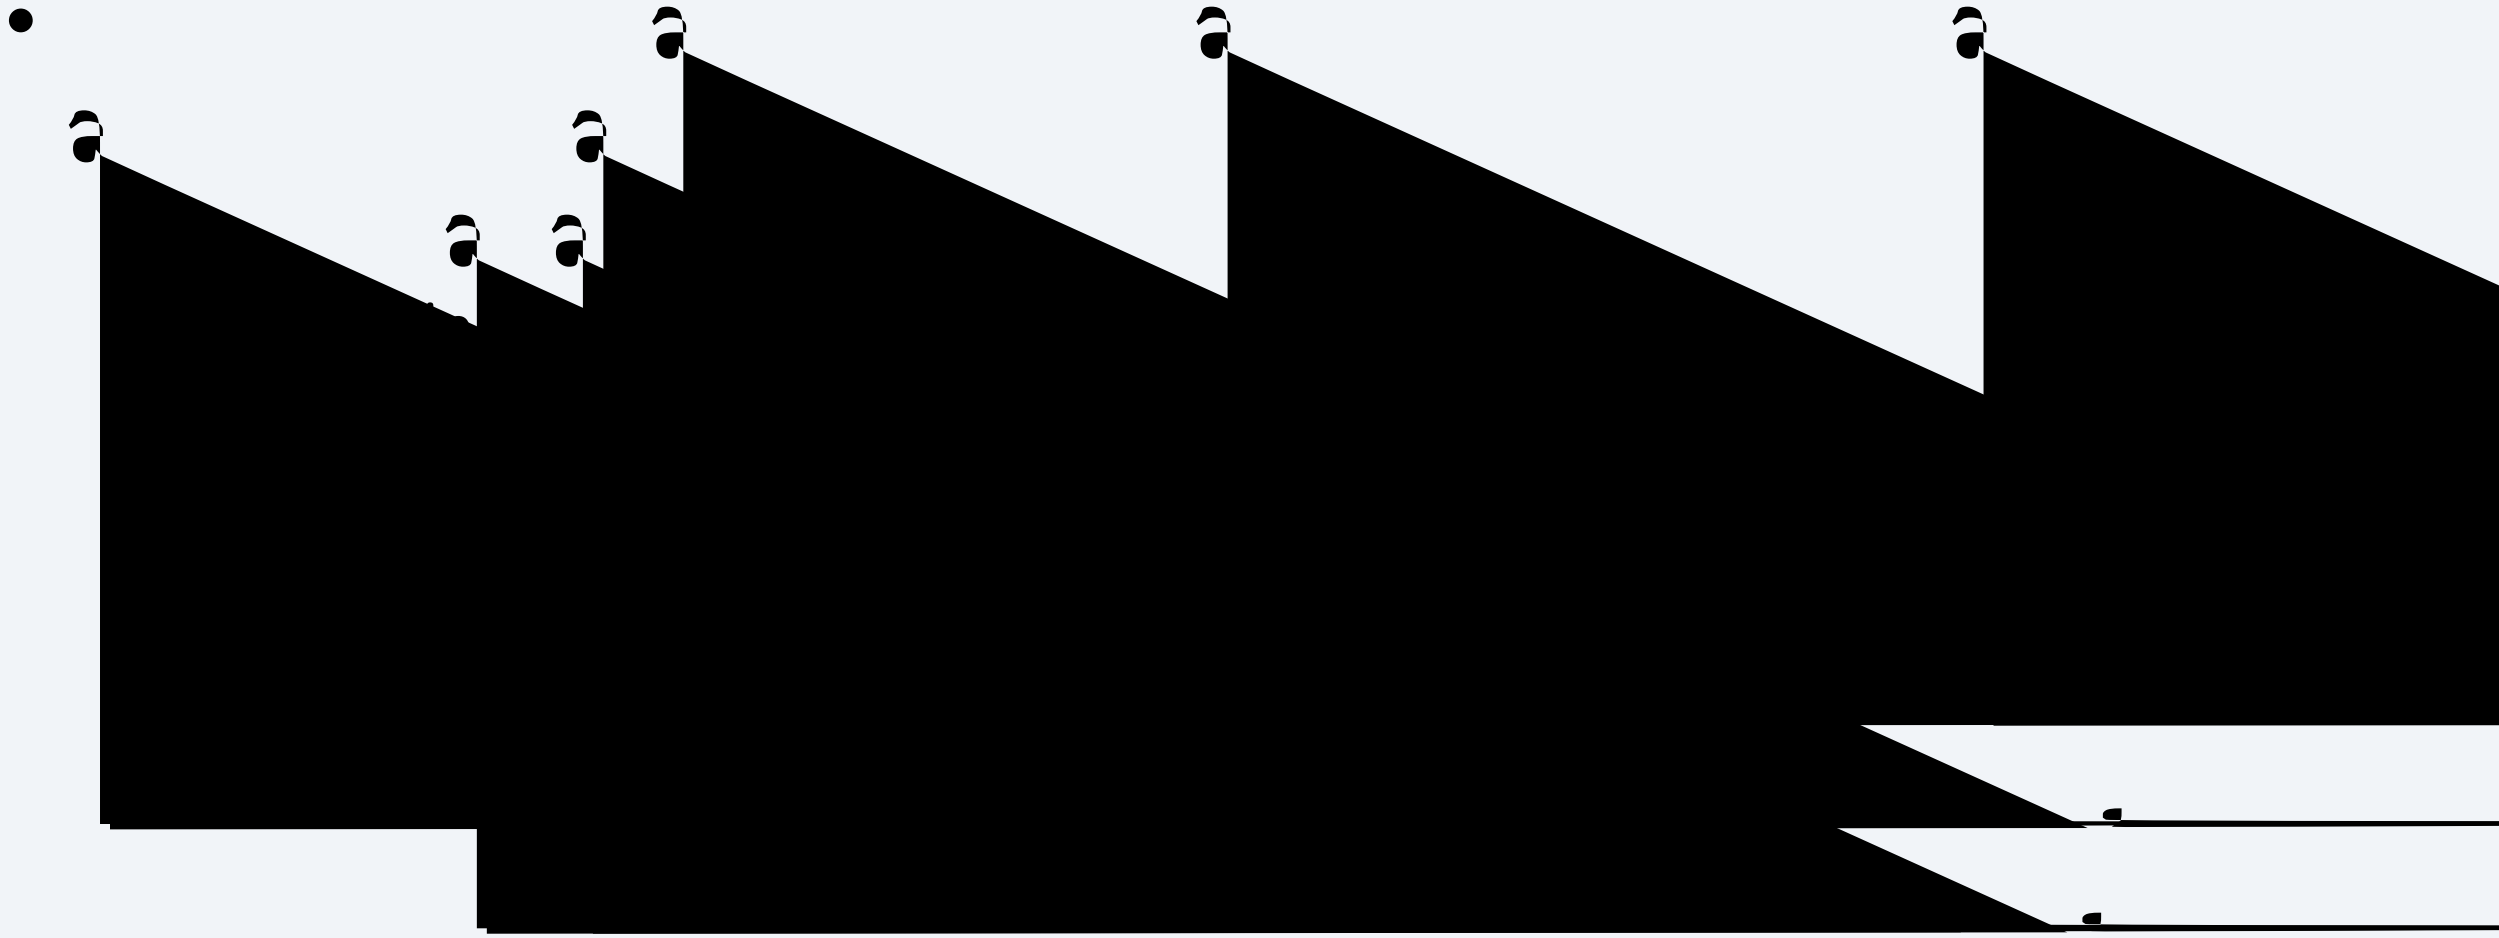

Problem 1: Single Channel = Finite Bandwidth

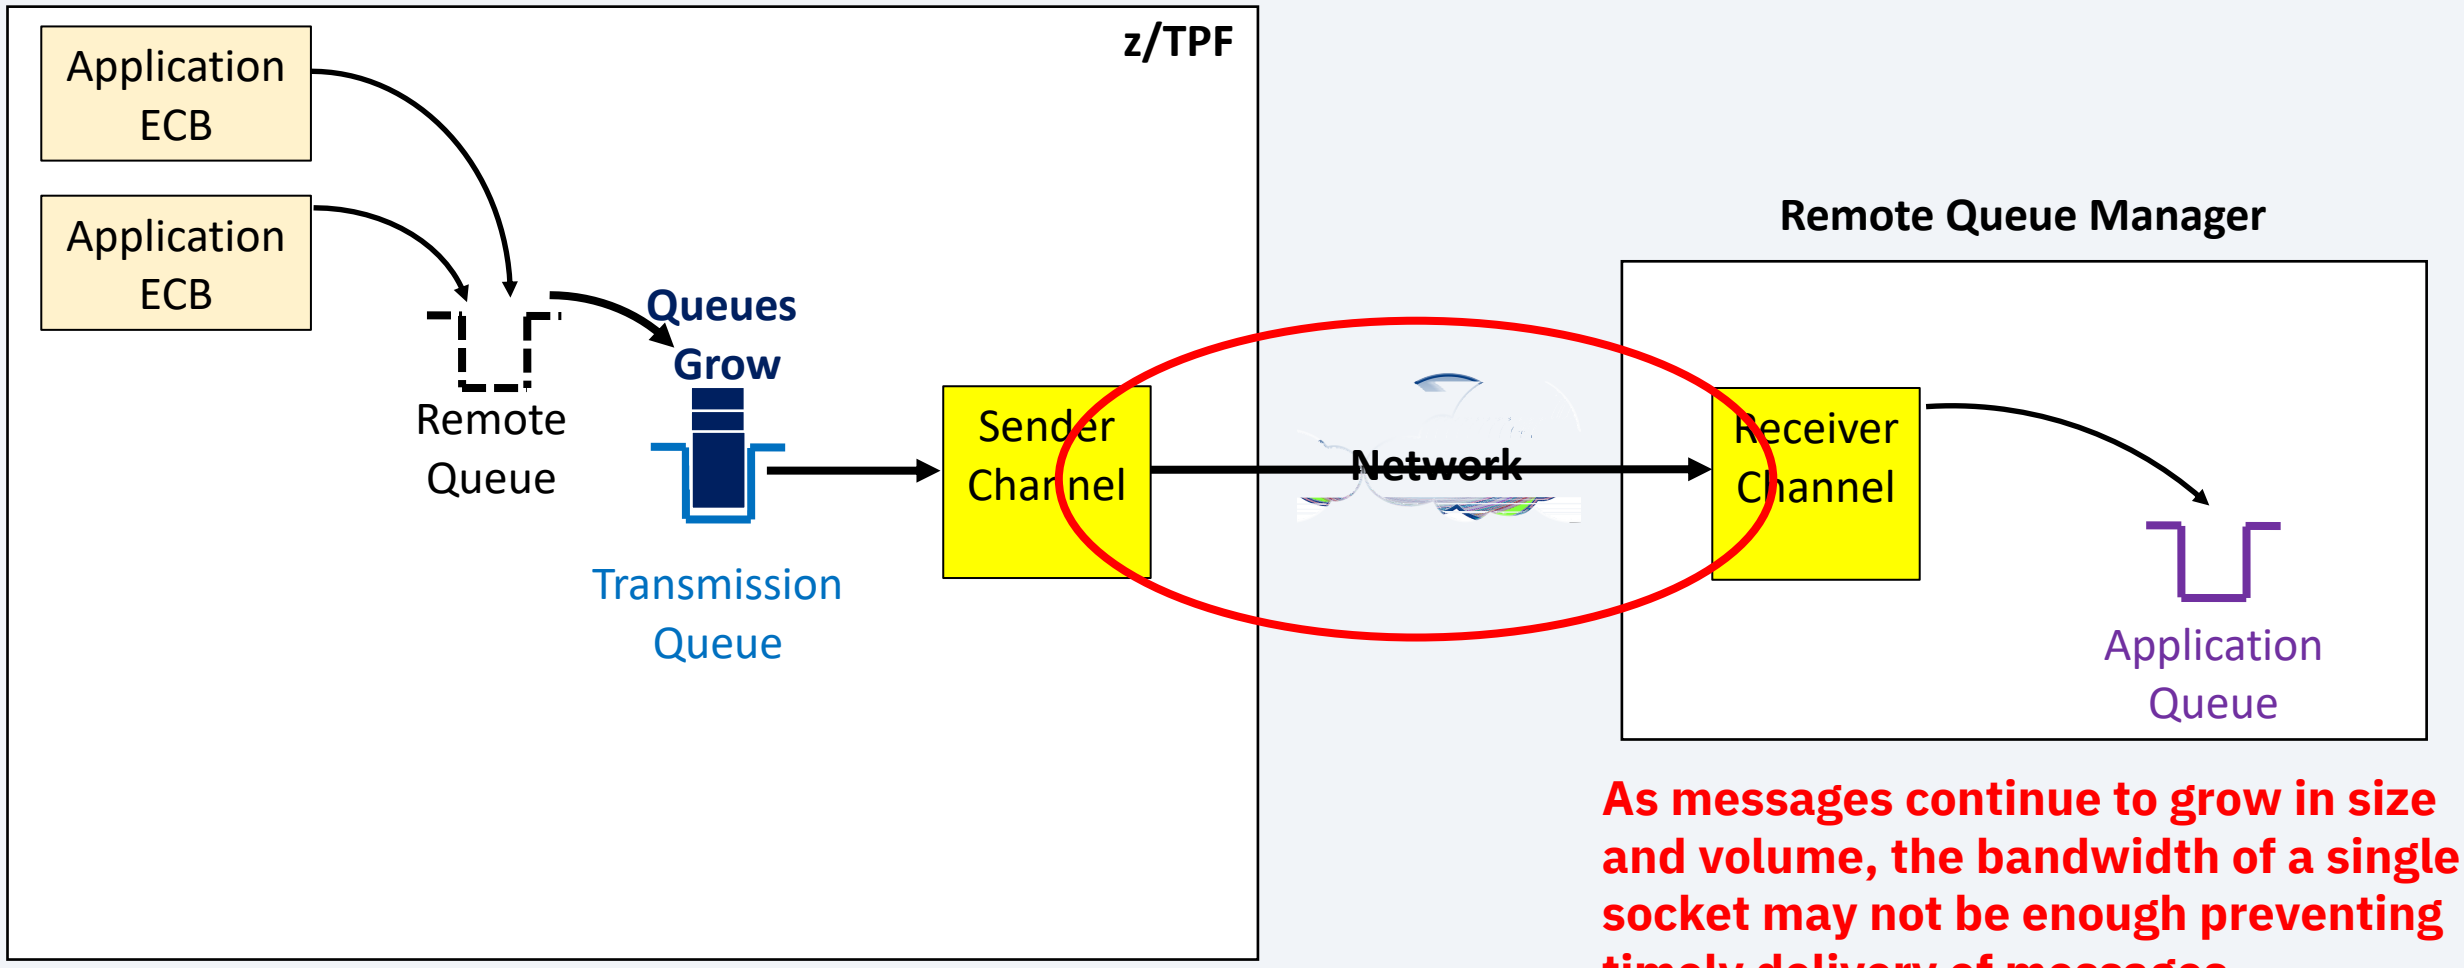

As messages continue to grow in size and volume, the bandwidth of a single socket may not be enough preventing timely delivery of messages.

Problem Statement

Problem 2: Overloading the Remote System

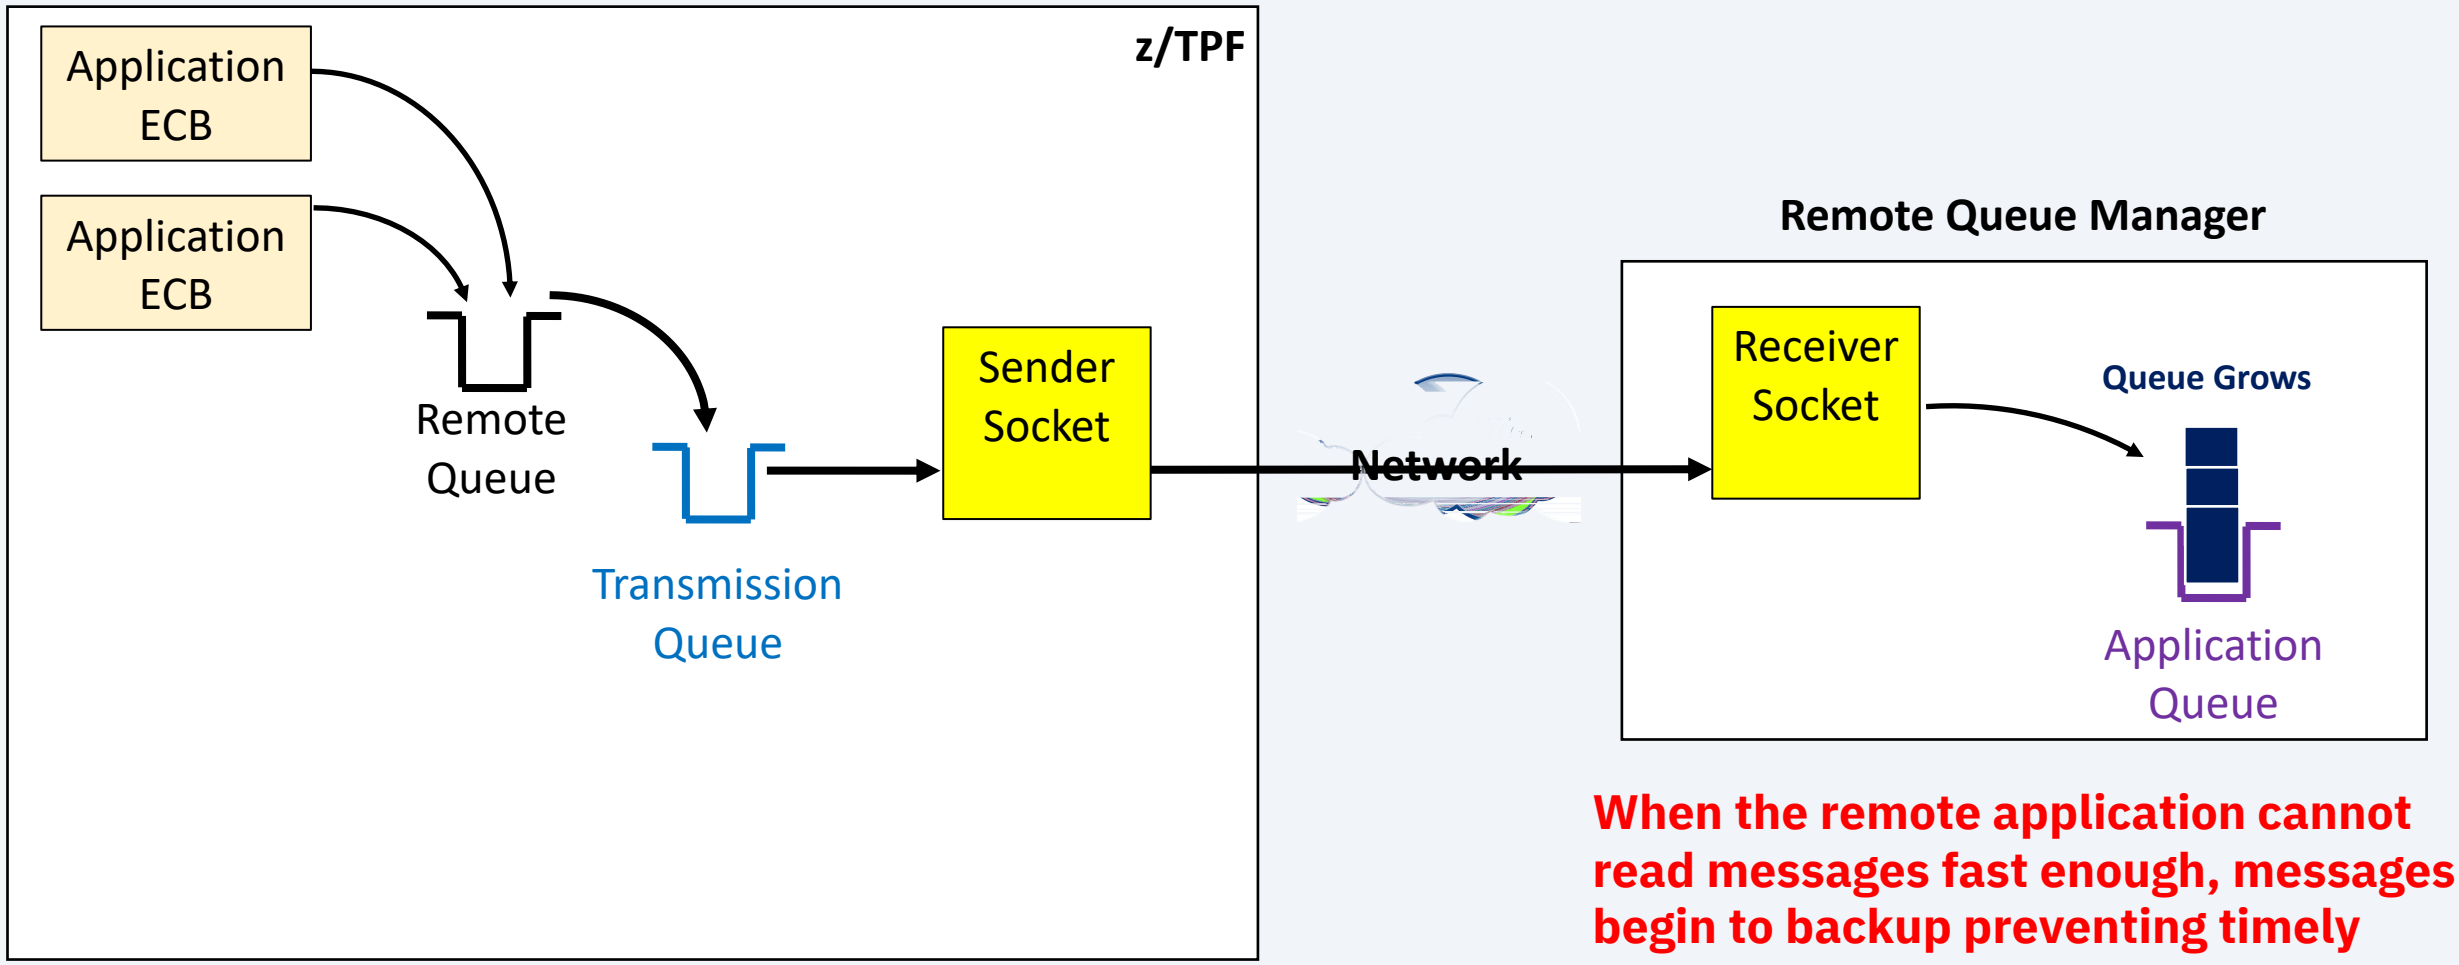

When the remote application cannot read messages fast enough, messages begin to backup preventing timely delivery of messages.

Problem Statement

Problem 3: Single Point of Failure

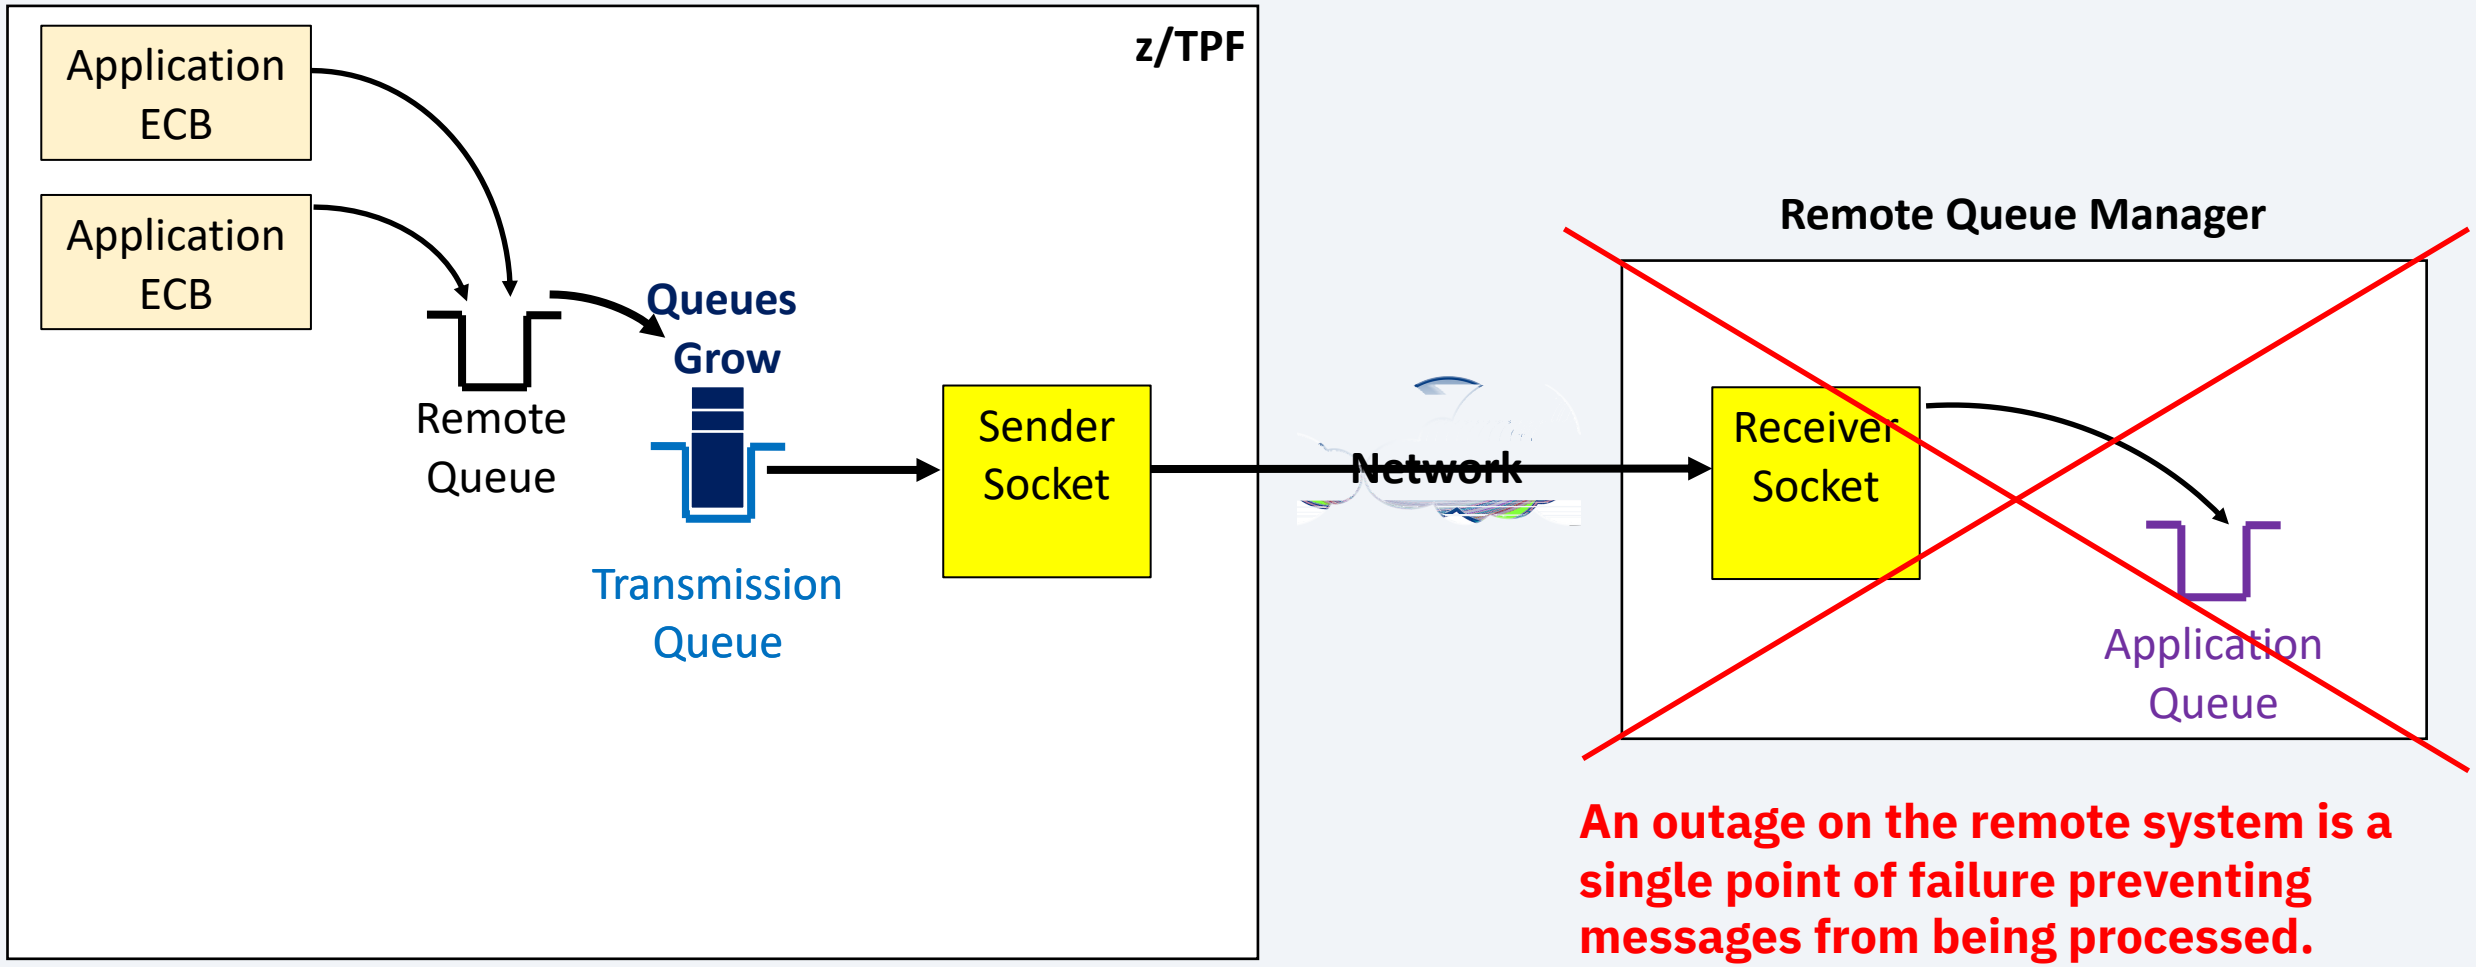

An outage on the remote system is a single point of failure preventing messages from being processed.

Creating Additional Transmission Queues

- **To prevent single point of failure or bandwidth constraints**
  - Additional channels can be created to the same system or separate systems

- **Applications must create code to monitor the status of channels and handle load balancing across them**
  - Increased application complexity to monitor the size of transmission queues
    - Access to the sender channel states is not available to the application
  - Adding new transmission queues may require application changes

Problem Statement

Application Managed Multi-Channels

Can create additional transmission queues application logic must be added

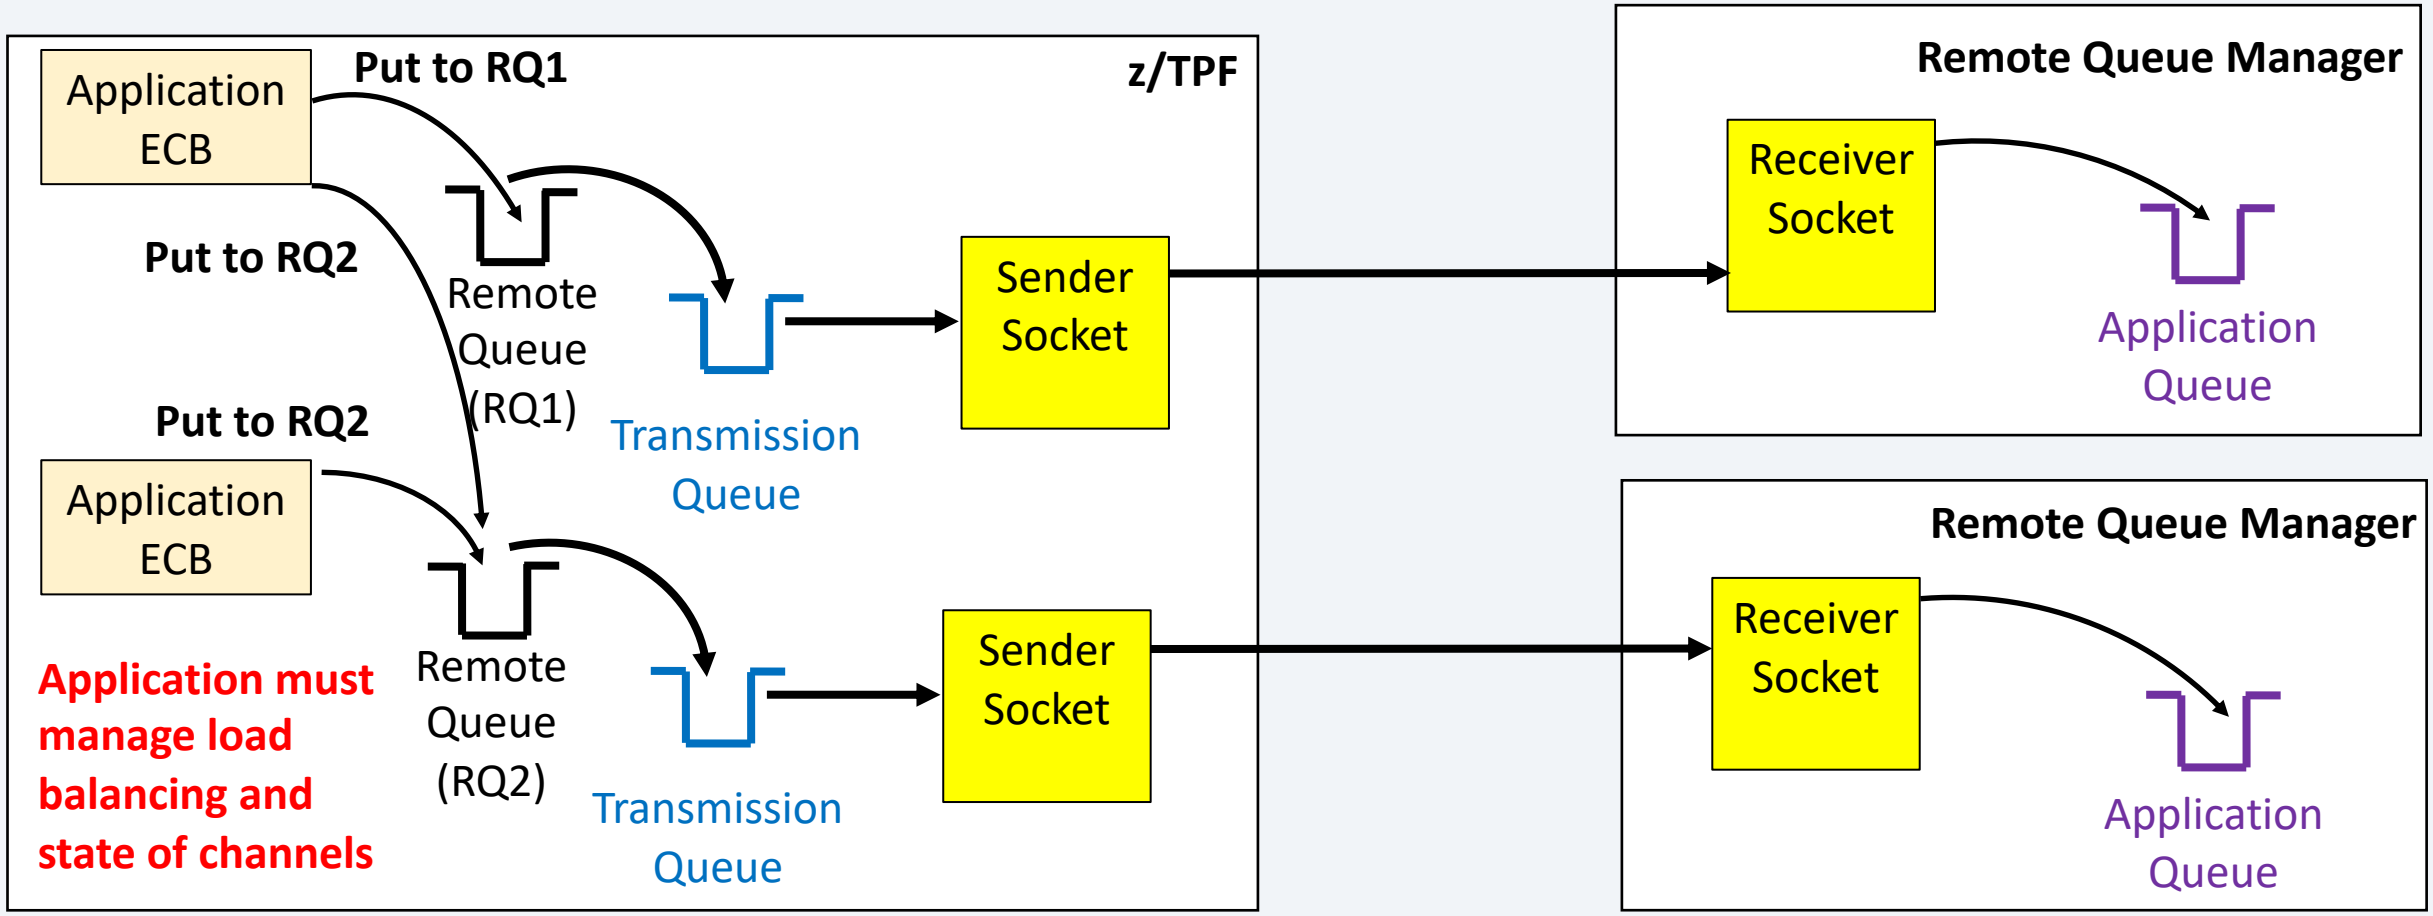

Sender channel sockets could be connected to the same, or a separate, remote MQ system

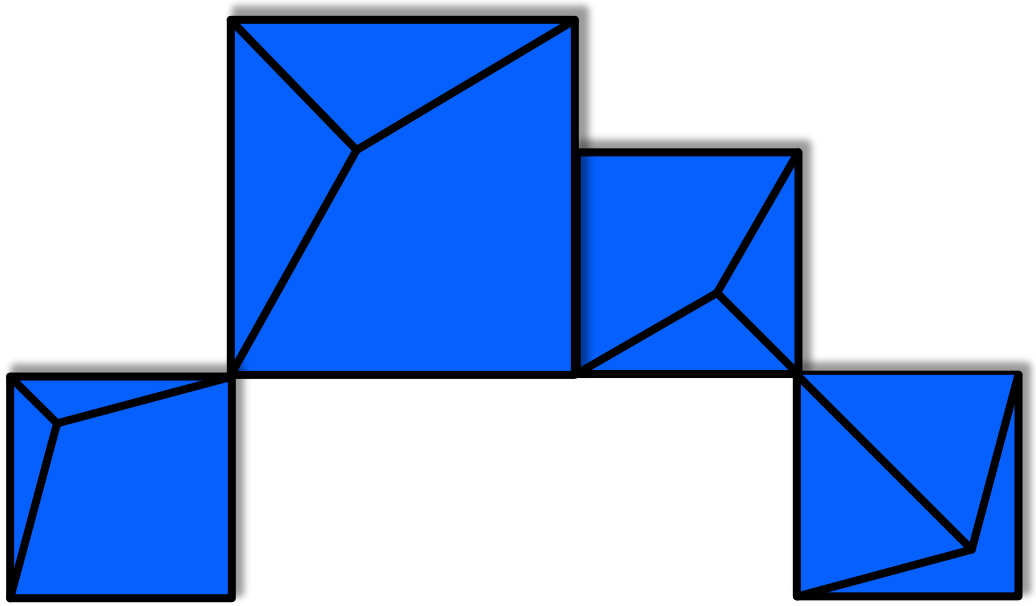

Solution

Solution

- Application always sends to a single remote queue name, but the remote queue name can be mapped to one or more transmission queues and sender channels.
  - Mapping is done administratively without impacting existing applications
  - Monitoring and management of the channels is handled by the z/TPF system

- Similar in concept to High Speed Connector where we administratively create groups of endpoints (remote systems).

Solution

z/TPF Dynamic Routing Architecture

Application ECB sending messages to a single remote queue name to invoke MQ Dynamic Routing

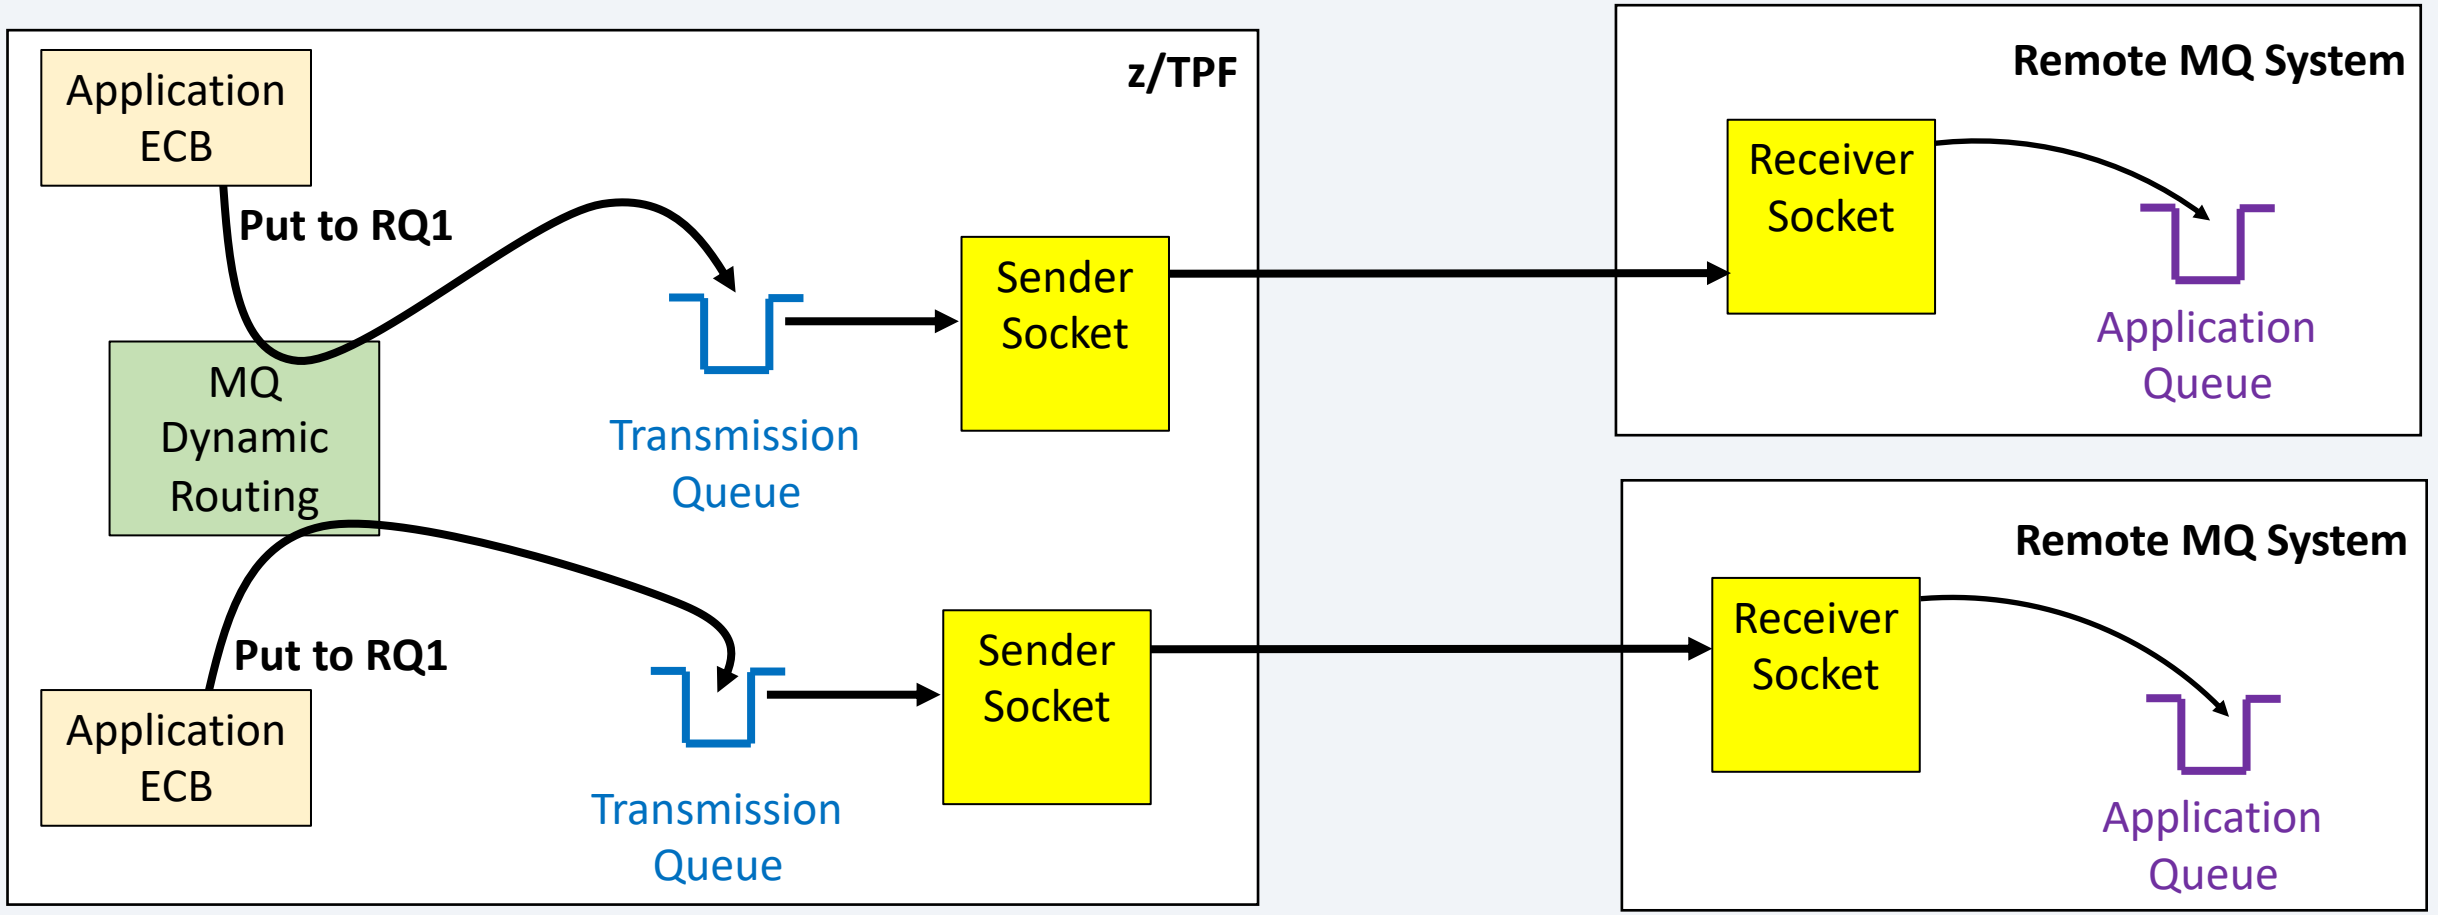

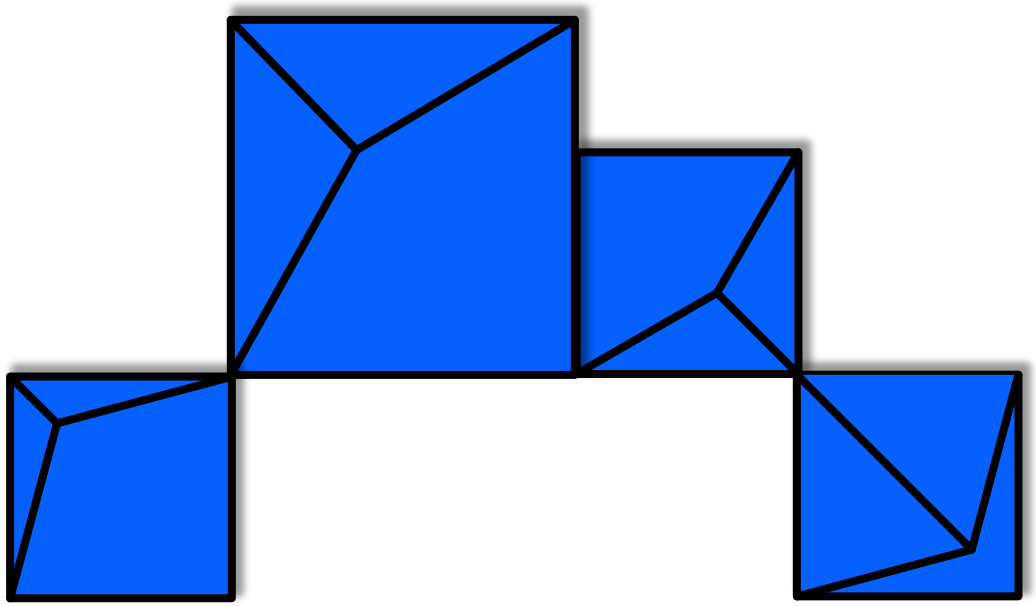

Technical Details

MQ Dynamic Routing Configuration File

MQ Remote Queue definitions are replaced with a new MQ Dynamic Routing configuration file.

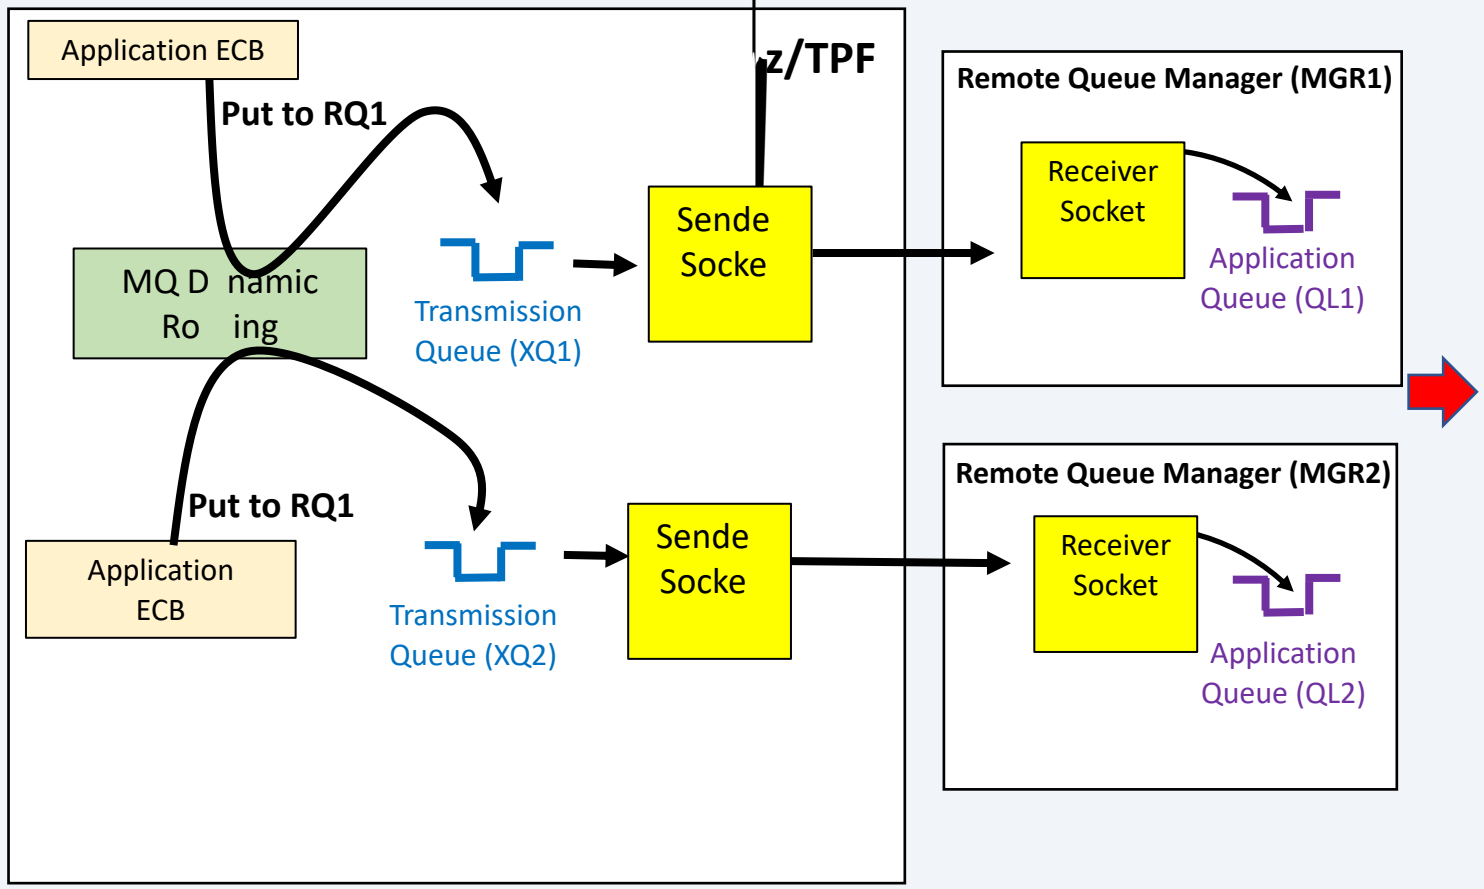

MQ Dynamic Routing Configuration File

```
RQ1,6,0,0,QL1,MGR1,XQ1
RQ1,6,0,0,QL2,MGR2,XQ2
```

MQ Dynamic Routing Configuration File

- C f g a f e aded e e f e e a da ed
  - T e f e e ga ed ce a ec d a d ef e ed f c a ged
- P ce e f e
  - e c d c c
  - e c d b c
- W S b c f e g
  - Rollout the same file across the complex
  - Keep backup versions of the files

MQ Dynamic Routing – Route Selection

- D a c R g de a a e c e a e MQOPEN e
  - O ce a e de e ed a b e e e age f ECB e a e a MQCLOSE ed
    - T ca be ef e e de f e age f a g e a ac a e
    - E e a e age f a a ca a ce d be e a e e e e e
  - If a ca e MQPUT e g be d e eac e e API ed
    - MQOPEN MQCLOSE d e de e c e b e e

MQ Dynamic Routing Migration

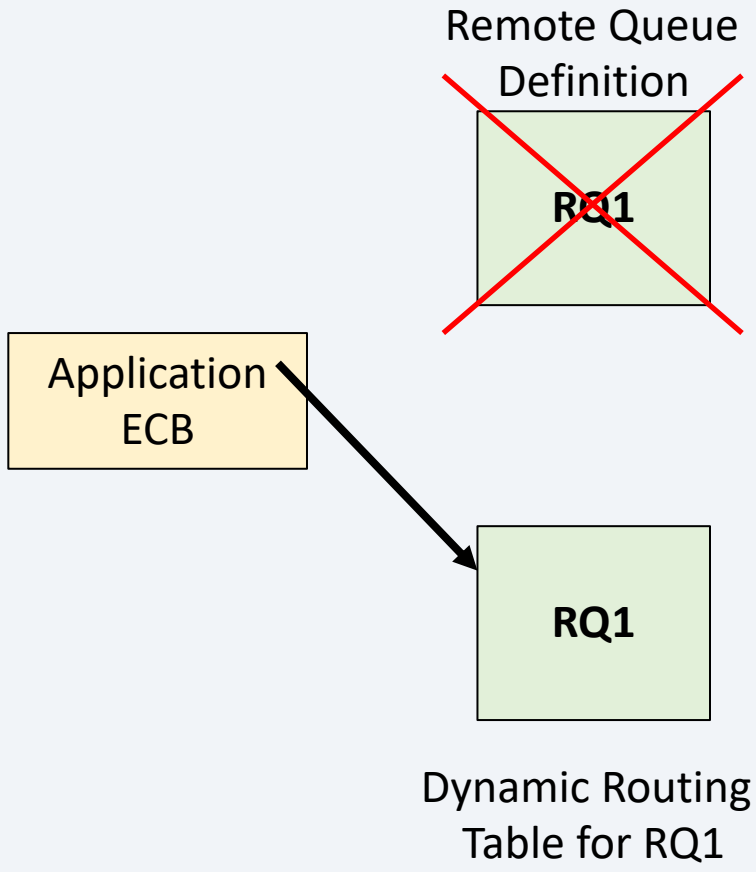

- When an application sends messages to RQ1
  - Remote queue definition is searched first
- After defining the dynamic routing table for RQ1 it remains inactive while the remote queue definition exists.
- Transitioning applications to use dynamic routing is done by deleting the remote queue definition
  - Fallback is to redefine the remote queue definition.

No application changes required!

MQ Dynamic Routing – Changing Topology

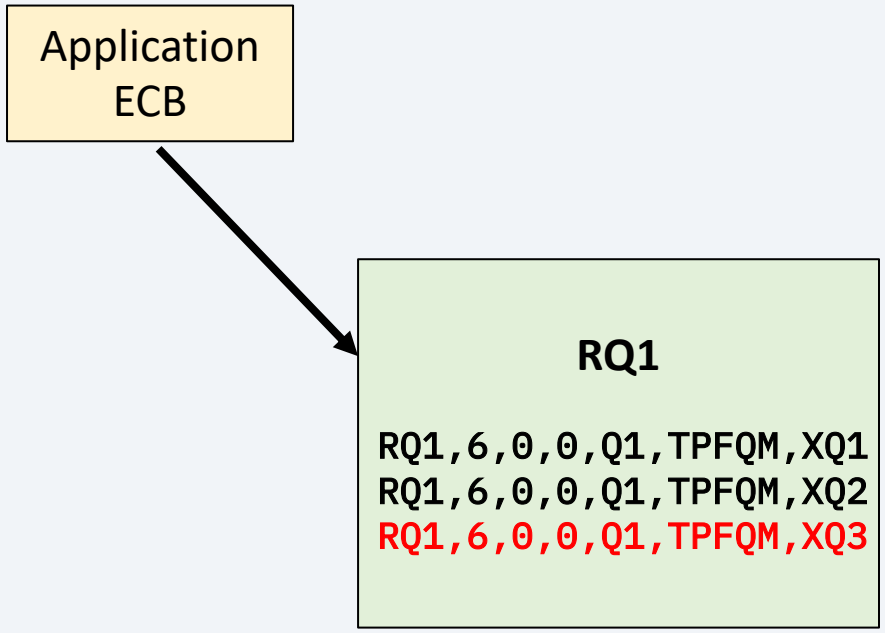

Dynamic Routing
Table for RQ1

- Easy to add additional channels to the MQ Dynamic Routing table.
  - Update and load the new file to z/TPF system
- z/TPF system will automatically detect and begin using the new channel definition.

No application changes required!

MQ Dynamic Routing Load Balancing

- D a c R e e e e e be c e b
  - T a e e a ac e e de c a e
    - S ed c a e be e a ed
    - U a a ab e e e a age be e a ed
  - L e e e de
    - Fa e e e a age ce e e age
  - W e a g a e e a
    - Q e e de a e ced e e ec ed

MQ Dynamic Routing Monitor and Management

- Ne ZMQDR c a d d a d a c g e e

```
User:     ZMQDR DISP QR-RQB1

System:   MMQDR0012I 13.51.52 DYNAMIC ROUTING SINGLE REMOTE QUEUE DISPLAY
          DESCRIPTION - Data Events to Remote Systems
          DEFPSIST    - NOT PERSISTENT
          PUT         - ENABLED
          ------------------------------------------------------------
          Transmission Queue Name                        Use Count    Avail
          ------------------------------------------------------------
          LNXQB1                                         - 1834113    - YES
             Remote Queue Name - LNXQ1
             Remote Qmgr Name - LNXMGR1
          LNXQB2                                         - 1839902    - YES
             Remote Queue Name - LNXQ1
             Remote Qmgr Name - LNXMGR2
          LNXQB3                                         -  830687    - NO
             Remote Queue Name - LNXQ2
             Remote Qmgr Name - LNXMGR3
          LNXQB4                                         - 1832723    - YES
             Remote Queue Name - LNXQ2
             Remote Qmgr Name - LNXMGR4
          END OF DISPLAY+
```

New ZMQDR RESET to reset the “Use Count”

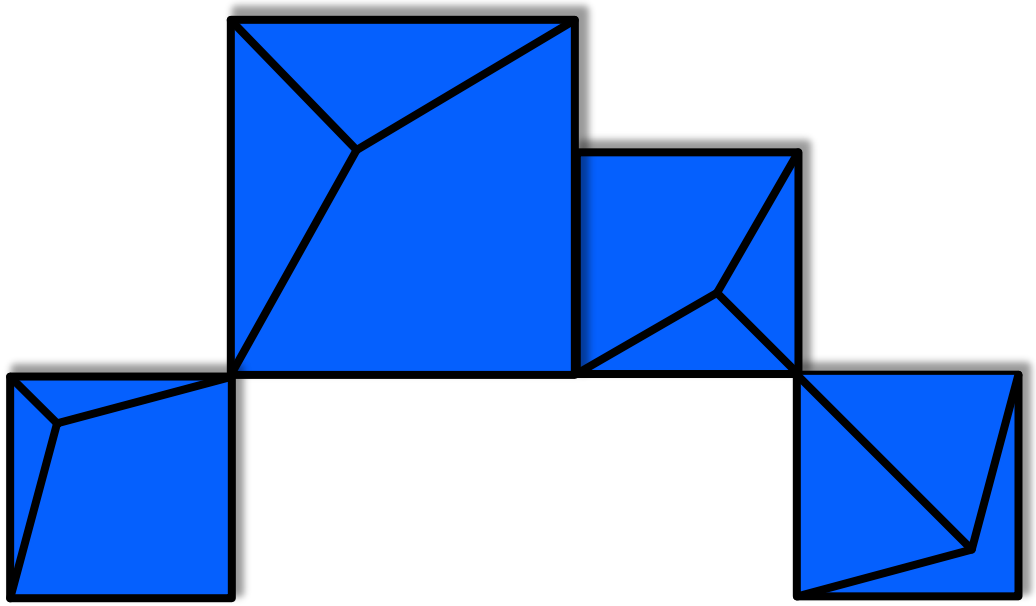

Value Statement

Value Statement

- **Delivered With APAR PJ45631 – November 2019**

- Multiple channels to remote systems that are managed and monitored by the z/TPF system
  - Increased throughput of messages being processed
  - Higher availability in the event of failures

- Intelligent load balancing across all available channels
- z/TPF managed routing based on the state of the system and the channels
- Does not require application changes or outages to spread application messages across multiple channels.
- Can dynamically add channel definitions to increase capacity

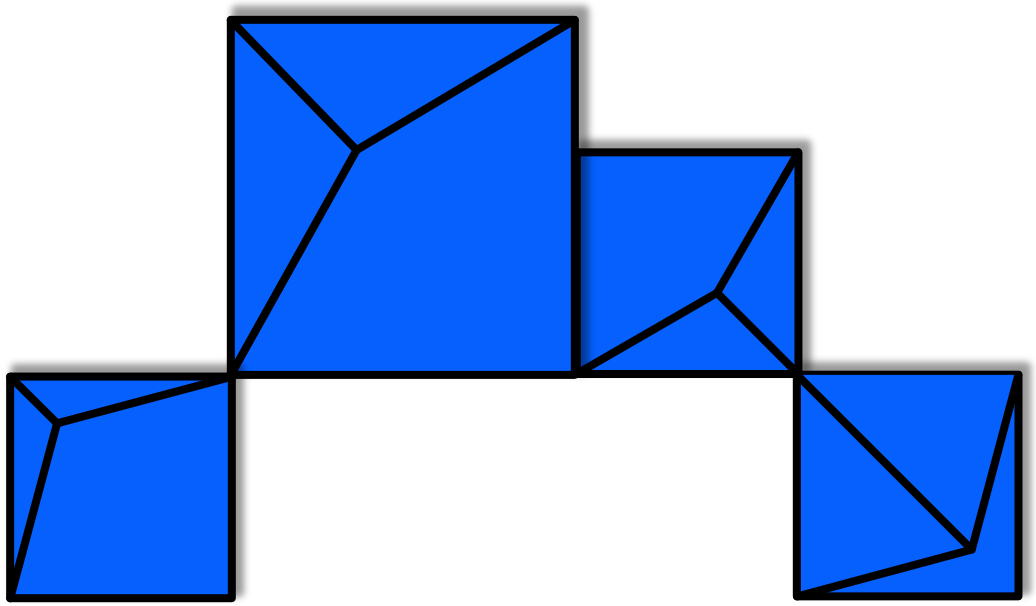

What’s Next

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.

Background

- Use of z/TPF MQ continues to increase.
  - For example, pushing data off of z/TPF with z/TPF Data Events

- z/TPF MQ uses System Work Block (SWBs) to save messages
  - SWBs are a critical system resource below the 2G bar
  - To minimize I/O, users define large amounts of SWBs if possible.

Pain Points

- Available space below the 2G bar can prevent future growth of the z/TPF system
  - For example, increasing the number of defined ECBs in the system
- With limited number of SWBs, more frequent I/O may result as the MQ sweeper needs to free up SWB storage
  - Small hiccups in the network or on the remote end
  - Surges of messages received on queues
- With increased usage of MQ and increase size of MQ messages
  - The network bandwidth consumed by MQ is growing
  - The network latency to send MQ messages to remote partners is increasing

Solution

A system administrator will see a reduction in I/O, CPU consumed, network costs, and processing time needed to process z/TPF MQ messages

- A system administrator will see a reduction in I/O, CPU consumed, and processing time used to process MQ messages by leveraging z15 HW compression when internally storing messages in memory or on disk.
- A network administrator will see a reduction in network bandwidth costs and reduced latency by MQ leveraging z15 HW compression support
- A system administrator will have the ability to define MQ queues that will use storage above the 2G bar to hold messages to enable workload growth

Technical Details

- New queue type that will use storage above the 2G bar
  - Primary use case is for high volume outbound data, mainly persistent MQ messages sent to remote nodes.

- Redesigned sweeper and checkpoint processing to reduce the number of I/Os and processing time
  - Leverage large memory to reduce I/O
  - Elimination of redundant I/O
  - Leverage z15 Hardware Compression to reduce I/O and CPU consumed

- New options to allow for z/TPF MQ network compression on channels

- Will be a phased deliverable – first being internal compression of MQ messages
  - Reduction of I/O and memory footprint for MQ.

Call for Sponsor Users

- Will be looking for Sponsor Users to assist in design and implementation, targeting the following personas:
  - z/TPF MQ administrators
  - z/TPF system administrators
  - z/TPF operators and coverage
  - z/TPF solution architects
- Targeting Sponsor User engagement to begin in September 2020
- For more information or to get involved as a sponsor user, contact
  - Jamie Farmer – jvfarmer@us.ibm.com
  - Danielle Tavella – Danielle.Tavella@ibm.com

Thank You

Questions? Comments?

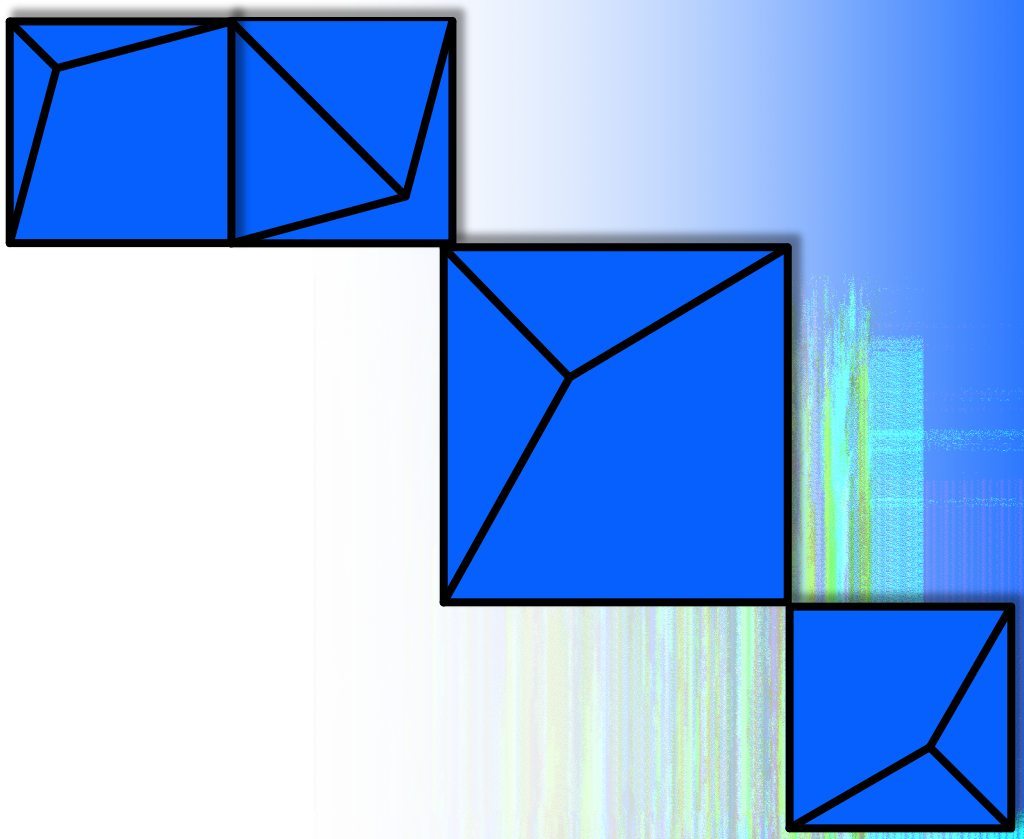

Trademarks

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.

Virtual TPFUG Q&A (Slide 1 of 2)

Summary of Q&A from the virtual TPFUG event:

Question	Answer
Q: So if updates for the same records went over differnt remote queue managers, it is up to the host TPF system to order them, correct? Disperate updates could be processed by differnt ECB's so the MQOPEN/CLOSE doesn't help.	A: Correct, the application must decide if it wants messages to go to the same destination. A single MQOPEN and multiple MQPUT's will send the messeges to same desitnition. If order and system does not matter, MQPUT1 can be used to route messages around.
Q: Would you anticipate that Dynamic Routing might eliminate the need to use existing SWINGQ feature?	A: For planned outages, when the transmission queue is disabled, dynamic routing can bleed the transmission queue to zero and negate the need for swingq. With an unplanned outage, the messages may still be on the tranmission queue, and SWINGQ is needed.
Q: Are there use cases that I should not prefer Dynamic MQ routing for?	A: If order is needed for messages, dynamic routing becomes problematic as messages may arrive out of order.
Q: Is that a similar concept than the cluster concept on WS MQ distributed side?	A: This is similiar to clustering, however does not support all of the clustering features. We choose to implement a smaller piece and this is z/TPF unique.
Q: with refresh process running every second, on a busy system, is there a risk of one process not completing in 1 sec, and having 2 processes trying to do the same process, blocking the Xqueue table??	A: The Dynamic Rotuing config file processing uses ENQ/DEQ to prevent two processes from updating the definition.
Q: Is dyanamic routing better than using active/passive failover MQ Qmgrs for resileincy ?	A: Dynamic Routing uses all available definitions and systems, I think active/passive is a standby system which this is not.
Q: Will this facilitate or provide the ability to move messages from one processor trans queue to another processor trans queue? Swing queue just moves messages from one queue to another queue on the same processor.	A: This does not route message from one processor to another. z/TPF MQ has MOVEMSGS if a processor is deactivated. TO2 queues (COMMON-YES) are sharted between processors, however Dynamic routing does not support them as transmission queues.
Q: would outbound compression happen before or after ip trace?	A: Compression is done at the middleware layer, which in this case in MQ, so that is before the TCP/IP layer, thus before IP trace A: MQ has it's own trace facility that can be used to view message data before it is compressed.

Virtual TPFUG Q&A (Slide 2 of 2)

Summary of Q&A from the virtual TPFUG event:

Question	Answer
Q: can u define a secondary qr that is only used when the primary set of Qrs are down. Or if it is up it will be in the rotation.	A: Currently, if its up it is in the rotation. There is no current concept of a primary / backup in dynamic routing...but does sound like it may be useful.
Q: John, I noticed in the XRemite Queue def you have a PUT Allowed = 0, can we setup 2 Remotes, one with PUT Allowed and one not, and only enable the second when the 1 is down? allows for backup, and since the same Queue is used, might save the data on the q?	A: The PUT allowed on the DR def, must be the same value, and is either enabled or disabled. What you want to do is disable the transmission queue with put disabled, that will work for what you want to do.