z/TPF V1.1

TPF Users Group – Fall 2012

z/TPF Communication & Cryptography Enhancements

Jamie Farmer
Communications Subcommittee

AIM Enterprise Platform Software
IBM z/Transaction Processing Facility Enterprise Edition 1.1.0

Agenda

- **TCP/IP enhancements**
- **Hardware cryptography enhancements**
- **Secure Socket Layer (SSL) enhancements**
- **Observations of using z/TPF public key infrastructure**

TCP/IP Send ECBs Causing Input List Shutdown Condition

- **In z/TPF, multiple ECBs can issue a send API on the same socket.**
  - System controls the serialization of these sends using ENQC/DEQC functions.
- **Send APIs complete immediately, unless the socket send buffer becomes full**
  - Remote not reading or not reading fast enough.
- **The number of ECBs waiting on a send could drive the system into input list shutdown even if send timeouts (SO_SNDTIMEO) are used.**

Enhancement to TCP/IP send() Processing to Detect Sockets with Multiple ECBs Waiting to Send – PJ40485

- **The z/TPF TCP/IP stack will now track the number of ECBs waiting on a send API for each socket.**
- **When the number of ECBs waiting to send on a socket reaches certain thresholds, the socket monitor user exit (USMO) user exit will be invoked.**
  - User exit invoked when queue length reaches 10, 25, 50, 100
  - User exit is invoked once for each threshold reached
  - Socket monitor resets the threshold.
- **Support does not affect the mainline path of send processing**

Socket Monitor User Exit (USMO) For ECBs Waiting to Send on a Socket

- **USMO is an existing user exit that is invoked when events occur on a socket**
  - For example, listener backlog exceeded
- **New event type for send queue ECB threshold reached**
  - IUSMO_SENDQ_LENGTH
- **Socket with the large queue is passed to the user exit**
  - Socket descriptor, local IP address, remote IP address, local port, remote Port
  - The send queue length for this socket is also returned
- **Allows user exit code to take action on this socket to prevent input list shutdown conditions.**
  - For example, the user exit can close the socket which will cause all the ECBs waiting to send to be posted.

ZSOCK DISPLAY to Display ECB Send Queue Length

```
ZSOCK DISPLAY FORMAT SOCK-C0000E
CSMP0097I 19.06.23 CPU-B SS-BSS  SSU-HPN  IS-01 _
SOCK0043I 19.06.23 TCP SOCKET CONTENTS FORMATTED
LOCAL IP -         9.057.013.251  LOCAL PORT  -            9999
REMOTE IP -        9.057.013.250  REMOTE PORT -            1074
PROTOCOL -               TCP   SOCKET TYPE -         STREAM   _
  .
  .
FRAGMENTS IN -             0  FRAGMENTS OUT -             0 _
SEND ECBS QUEUE THRESHOLD -   100  SEND ECBS QUEUE LENGTH -      67
CLOSE ISSUED -             N
DNS NAME - tpfosa2h122.pok.ibm.com
AOR PENDING -              N
AOR TOKEN -                AOR PROGRAM NAME -
SOCKET CREATED - MON OCT 08 19.05.39 2012
END OF DISPLAY
```

z/TPF select API and OpenLDAP

- **Socket select API allows for an application to monitor sockets for readability and writability**
- **For example, is there data available to read**
  - Can pass an array of socket descriptors to select on
- **OpenLDAP issues select on an array of descriptors**
  - Listener socket
  - Client connected sockets for this LDAP server

z/TPF select API Issue Affecting OpenLDAP – PJ39769

- **z/TPF select returns an incorrect return code when a socket being selected is abnormally terminated.**
  - For example, RST received on a connected socket.
  - In this case, the select API returns a error return code with a SOCNOTSOCK error number.
- **Architecturally, the select API should return a good return code**
  - Socket that was abnormally terminated should be marked ready
- **Causes problems with OpenLDAP as a negative return code on select is considered a logic error**
  - Causes the LDAP listener to end in error.
- **The z/TPF Select API was updated to conform to standard**
  - This could have an affect on existing applications if selecting on more than one socket descriptor.

Hardware Acceleration For Public Key Cryptography

- **Public key cryptography requires a significant amount of overhead if operations are performed in software.**
  - RSA encryption / decryption
  - Used during SSL session startup
- **Prior to PUT 9, the z/TPF system supports the following hardware accelerators for public key cryptography**
  - PCI Cryptographic accelerator (PCICA)
  - CryptoExpress2 (CEX2A)
  - CryptoExpress3 (CEX3A)

Crypto Express4S Support – PJ40362

- **Crypto Express4S is the new generation of cryptographic hardware for zEnterprise EC12**
  - As in previous generations, the CryptoExpress4S can run in different modes
- **z/TPF now supports the Crypto Express4S running in accelerator mode (CEX4A)**
  - Accelerator mode is optimized for performing RSA operations

Using the tpf_cryptc() API to Encrypt and Decrypt Data

- **The tpf_cryptc() API provides a means for an application to encrypt and decrypt data**
  - Key is passed by the application (Clear Key)
- **The tpf_cryptc() API supports the following cipher algorithms**
  - DES, TDES, AES128, and AES256
- **The API will perform the operation in hardware if Central Processor Assist for Cryptographic Functions (CPACF) hardware exists for DES, TDES, AES128**
  - Operation performed in software (openSSL) if hardware does not exist
- **Prior to PUT 9, the AES256 cipher operation is performed in software only.**

AES-256 hardware acceleration for the tpf_cryptc() clear key API – PJ40018

- **The tpf_cryptc() API now supports the AES256 cipher algorithm to be performed in hardware.**
  - Operation will be performed in CPACF hardware when available.
  - If hardware is not available, operation is performed in software.

Stopping Shared SSL Sessions

- **When a ZSSLD STOP command is issued, all shared SSL daemons are stopped**
  - The SSL sessions allocated to each daemon are cleaned up
- **If SSL sessions are cleaned up that have an ssl_aor pending, a new ECB is created for each of those SSL sessions**
- **Depletion of ECBs can occur on ZSSLD STOP if enough SSL sessions have ssl_aor pending.**
  - SESSsec parameter could be used on the command to throttle the SSL session cleanup
    - If parameter is left off or value specified for it is too high, ECB depletion can occur

ZSSLD STOP Enhancement to Prevent ECB Depletion – PJ39830

- **ZSSLD STOP processing enhanced to prevent ECB depletion and possible CTL-064C04 system error.**
- **Processing now tracks the number of SSL sessions with ssl_aor pending that have been cleaned up.**
  - Periodically issues DLAYC to throttle new ECBs for SSL sessions with ssl_aor pending.
- **This processing will take effect regardless whether the SESSsec parameter is specified on the command.**

Shared SSL Performance Issue

- **For every shared SSL API call, the shared SSL code must clear the daemon's error queue**
  - The SSL code acquires three locks to perform this functionality
  - The locks obtained are shared across all threads on all daemons – increasing lock contention
  - Locks are performed using z/TPF CORHC which can cause performance degradation when the number of ECBs waiting increases.

Shared SSL Performance Enhancements – PJ39057, PJ39830

- **Processing has been enhanced to eliminate the bottleneck when clearing the shared SSL thread's error queue.**
  - Reduced the number of locks needed from three down to one–PJ39057
  - Different lock is obtained for each SSL daemon process. This decreases the amount of lock contention – PJ39057
  - Clearing error queue now uses LOCKC as opposed to CORHC to reduce ECB buildup and improve performance – PJ39830

Using z/TPF Secure Public Key Infrastructure – Cryptographic Accelerators Are Required

- **Using public/private key pairs saved in the z/TPF Secure Keystore requires cryptographic hardware accelerators**
  - For example, the Crypto Express3 accelerator (CEX3A)
  - Hardware performs expensive RSA encryption / decryption
  - Performing operations in software is not secure
- **On more than one occasion, customers have tried to use Secure Public Key Infrastructure without the required hardware**
  - For example, starting SSL sessions using public/private keys from the secure keystore
  - Without the hardware, SSL sessions fail and it isn't obvious the reason for failure
- **Warning messages have been added as part of Crypto Express4S enhancement (PJ40362)**
  - At key generation time if no cryptographic hardware is available
  - At IPL time if any secure RSA keys exist but no cryptographic hardware is available

Summary

- **Socket monitor enhancement to detect many ECBs queued waiting to send on a socket – PJ40485**
- **Change to z/TPF select processing to prevent issues with OpenLDAP and potentially customer applications – PJ39769**
- **z/TPF support for Crypto Express4S** (zEnterprise EC12) **– PJ40362**
- **Enhance tpf_cryptc() clear key encryption API to allow AES-256 operations to be performed in hardware – PJ40018**
- **Enhance shared SSL support to throttle SSL session cleanup when stopping the shared SSL daemons – PJ39830**
- **Shared SSL Performance Enhancement – PJ39057, PJ39830**

Trademarks

- **IBM, the IBM logo, and ibm.com 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.**