z/TPF Communication and Security Enhancements

Raymond Fan
z/TPF Development

Agenda

- TCP/IP Performance Improvements
- Increasing the number of OSA read buffers
- z/TPF hardware generated random numbers
- 512-bit SHA-2 Message Digests
- CPACF Performance Improvements

TCP/IP Performance Improvements

TCP/IP Hash Table Background

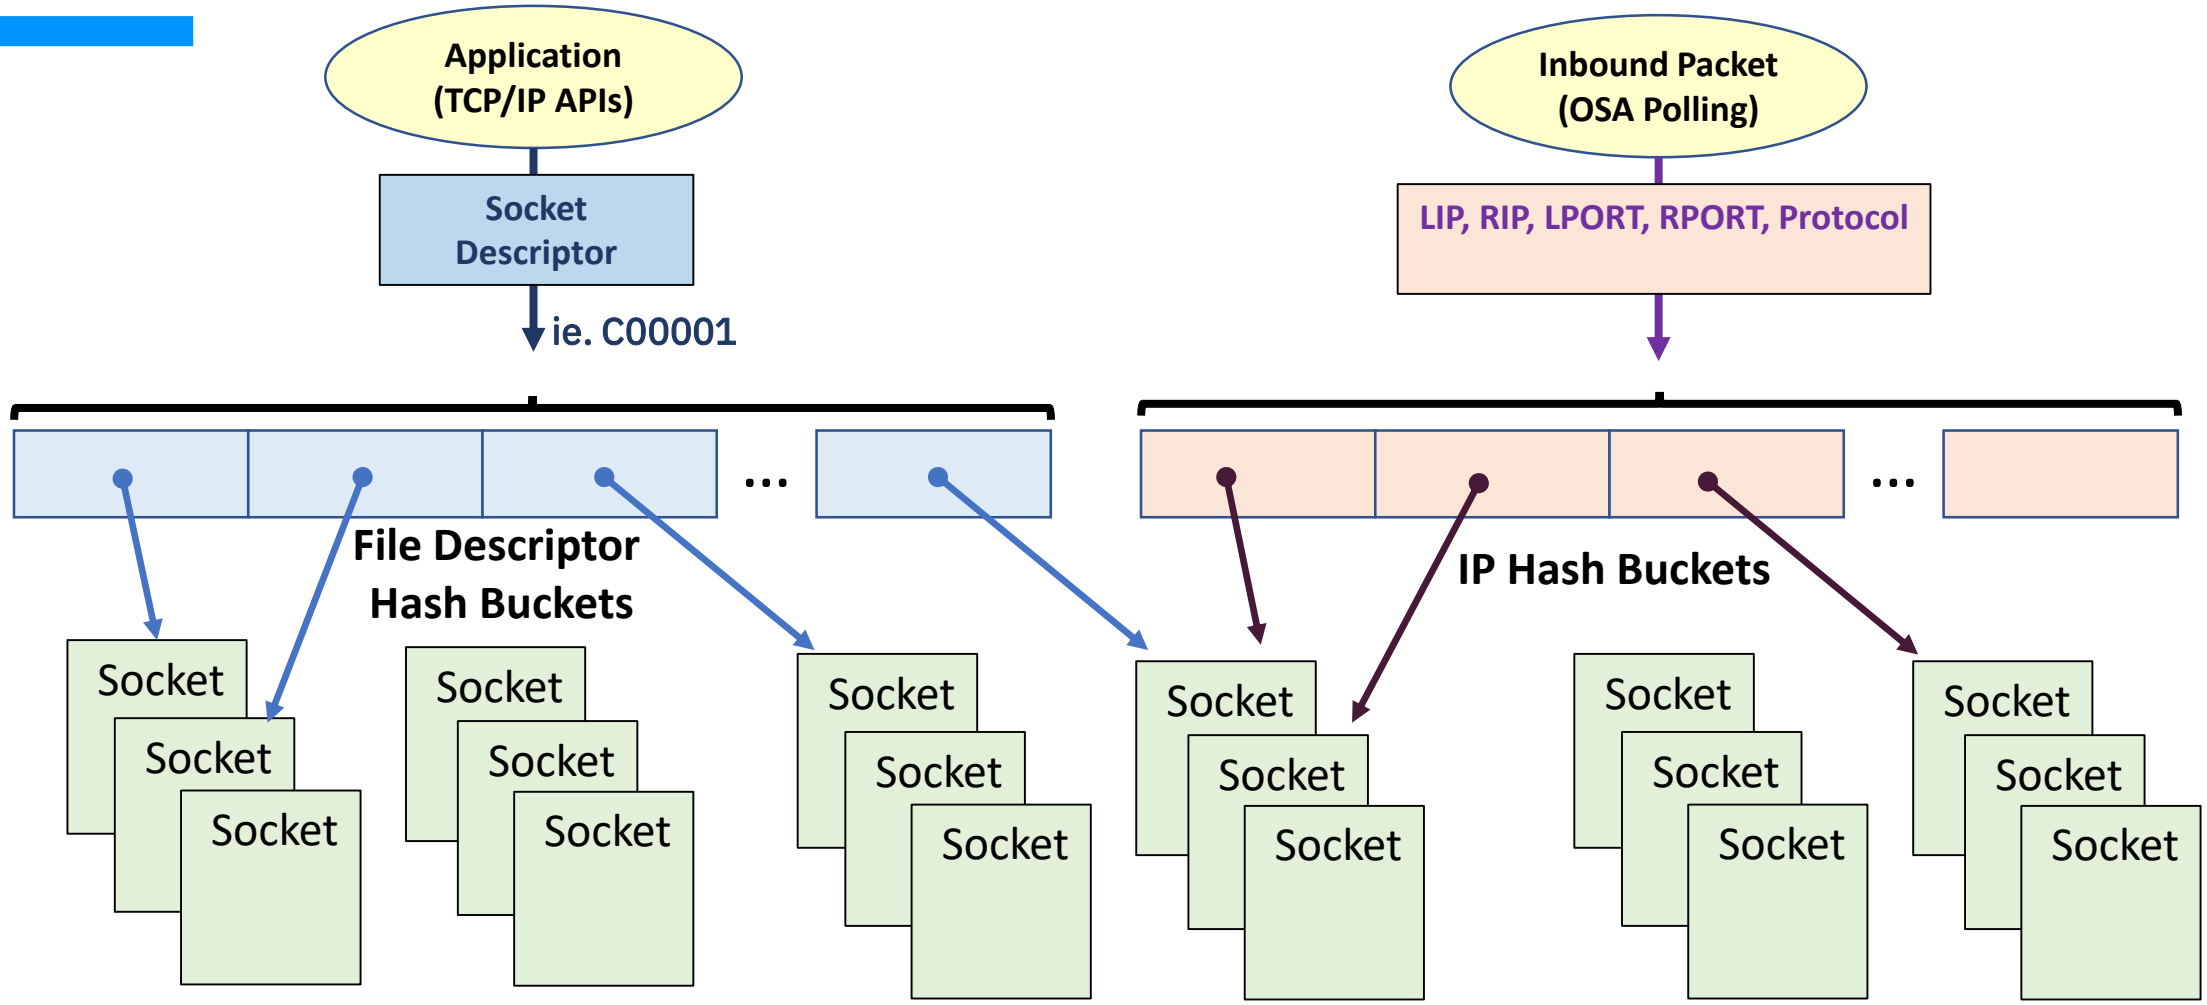

TCP/IP Hash Table Problem

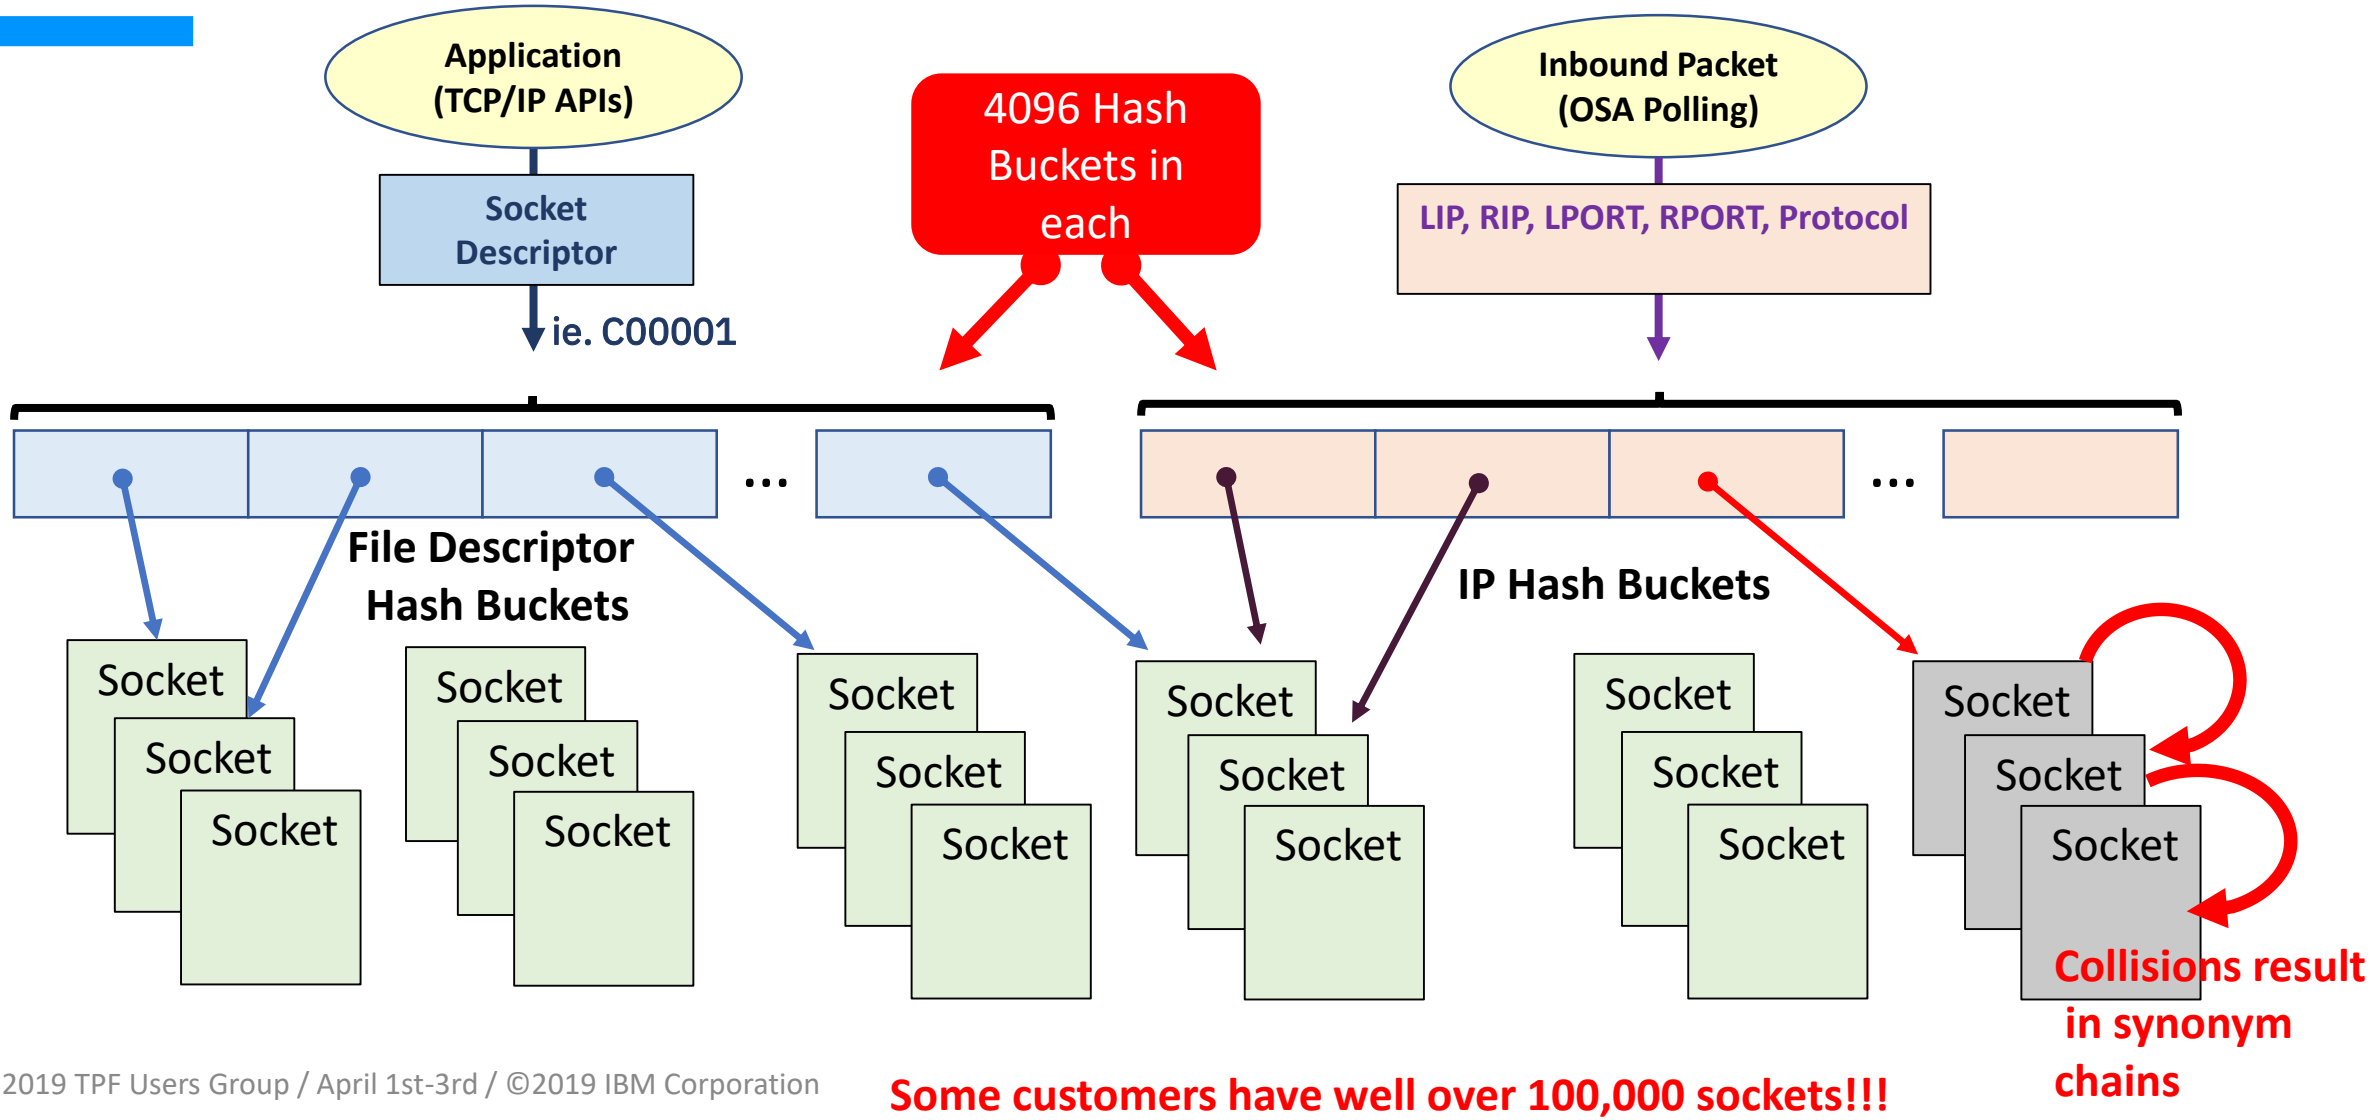

Some customers have well over 100,000 sockets!!!

Socket Hash Bucket Statistics

- With 4096 hash buckets and 100,000 sockets
  - On average each hash bucket will have a synonym chain of 24 entries
  - On average, each hash bucket lookup will scan 12 sockets before finding the target.

Increasing the Socket Hash Buckets

- TCP/IP hash buckets have been increased to 128K (131,071)
- Performance improvement depends on the number of active sockets
  - Up to 32x reduction in synonym chain overhead as number of active sockets approaches 1 million
- Delivered 3Q 2018 – APAR PJ45093

Increasing OSA Read Buffers

OSA Read Buffer Processing

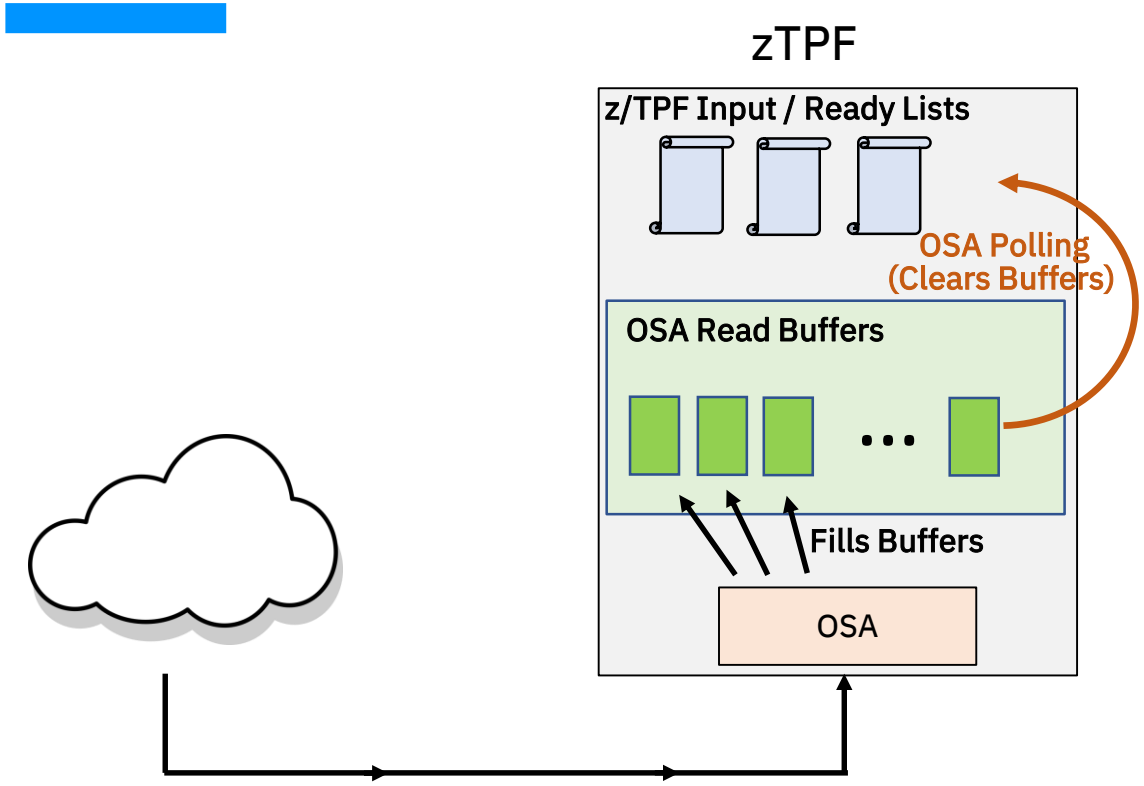

- Configurable number of OSA read buffers per OSA Connection
  - 16, 32, 64
  - Each buffer can hold up to 64K of inbound data

OSA Read Buffer Full Conditions

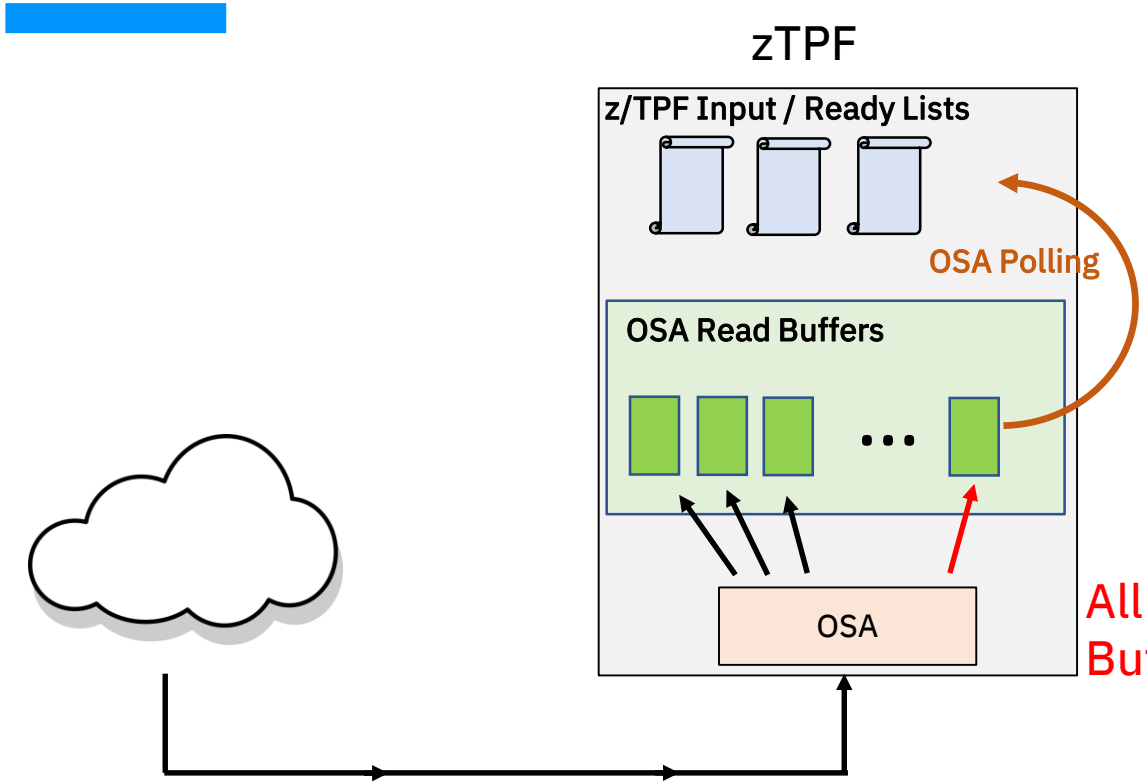

- When polling cannot be called, all the inbound buffers can fill
  - Buffer overrun
- OSA card will queue some inbound data, but eventually inbound packets will be dropped.

```
TTCP0323W 11.33.50 ALL READ BUFFERS FULL FOR OSA-OSASHAP
                   PACKETS MIGHT BE LOST
                   MILLISECONDS SINCE OSA POLLING WAS CALLED - 18 MS
```

Increasing the OSA Read Buffers

- Increasing the number of OSA read buffers to the architected maximum of 128 buffers.
  - No longer a configurable value in keypoint 2
- Delivered 1Q 2019 – APAR PJ45555
- In addition, improvements made to reduce the time between OSA polling and ensure it is called in a timely manner.
- These changes will minimize the number of buffer overruns that occur on the z/TPF system.

OSA Buffer Display Changes

```
ZOSAE BUFFER OSA-OSA1
CSMP0097I 14.33.33 CPU-C SS-BSS  SSU-HPN  IS-01
OSAE0013I 14.33.33 READ BUFFER USAGE FOR OSA-OSA1
 BUFFERS     INSTANCES      BUFFERS     INSTANCES      TIME
 -------  -------------   -------  -------------   ----
    0       132 185 100     16-31           7 217      5
    1        79 380 026     32-47               0      0
    2        14 351 998     48-63               0      0
    3         4 698 873     64-79               0      0
    4         2 337 560     80-95               0      0
    5         1 405 513     96-111              0      0
    6           950 074    112-125              0      0
    7           672 021
    8           472 965
    9           329 062
   10           230 114
   11           168 469
   12           132 836
   13           110 890
   14            97 051
   15            85 908

 Number of buffer full conditions: 0
 Number of full conditions caused by system error: 0
 Average time between polling during buffer full condition: 0
 TOD of last buffer full condition: N/A
 TOD of last buffer statistics reset: D5A729B125C22854
END OF DISPLAY
```

- Greater than 16 buffers in use is separated into buckets
- Average time between polling calls is displayed

TCP/IP Hardware Generated Random Numbers

Hardware Generated Random Numbers

- IBM EC12
  - Introduced an assembler instruction to create deterministic random numbers in hardware
  - Users of the DRNG are required to supply the seed

- IBM z14
  - True Random Number Generation (TRNG) using the same DRNG hardware instruction.
  - TRNG does not require the generator to be seeded.

Random Number Generation is done within the CPACF coprocessor.
The same hardware that does message digests and encryption.

Creating Hardware Generated Random Numbers on z/TPF

- A new tpf_random() function has been created to create hardware generated random numbers

```
int rc = tpf_random(rand_addr, rand_size);
```

- Where
  - rand_addr is the address of where to place random data
  - rand_size is the length of random data to create

The tpf_random() Function Details

- The tpf_random() function will use the deterministic random number generator (DRNG) to create the random number
  - The true random number generator (TRNG) is only used to supply the seed to the deterministic random number generator
    - Creation of true random numbers (TRNG) is expensive

- The seeding and reseeding of the DRNG is handled internally by the z/TPF system

When Does the z/TPF System Reseed

- The z/TPF system will reseed the deterministic random number generator
  - When the number of tpf_random requests exceeds the configurable value in keypoint 2 (CTK2)
  - When a tpf_random request has not been called for more than 10 seconds

- Using ZNKEY, an administrator can immediately change how often reseeding takes place (no IPL required)
  - ZNKEY RANDSEED-4096
    - 4096 tpf_random calls before reseeding (default is 1024)

Hardware Generated Random Numbers and Processor Levels

- If on IBM z14 or higher
  - The tpf_random function is fully supported.

- If on IBM EC12 or IBM z13
  - The tpf_random function is supported, but a user exit (URND) must be coded to supply the seed.

- If below the IBM EC12
  - The tpf_random function is not available for use

Querying the DRNG/TRNG Capabilities

- Use the ZCPAC QUERY command to determine what is supported on the processor the z/TPF system is running on:

```
ZCPAC QUERY

CSMP0097I 10.45.26 CPU-A SS-BSS  SSU-BSS  IS-01
CPAC0012I 10.45.26 CPACF QUERY DISPLAY
   SHA-1:        ENABLED
   DES/TDES:     ENABLED
   AES-128:      ENABLED
   SHA-256:      ENABLED
   AES-256:      ENABLED
   SHA-512:      ENABLED
   DRNG:         ENABLED
   TRNG:         ENABLED
```

z/TPF Random Number Generator Architecture

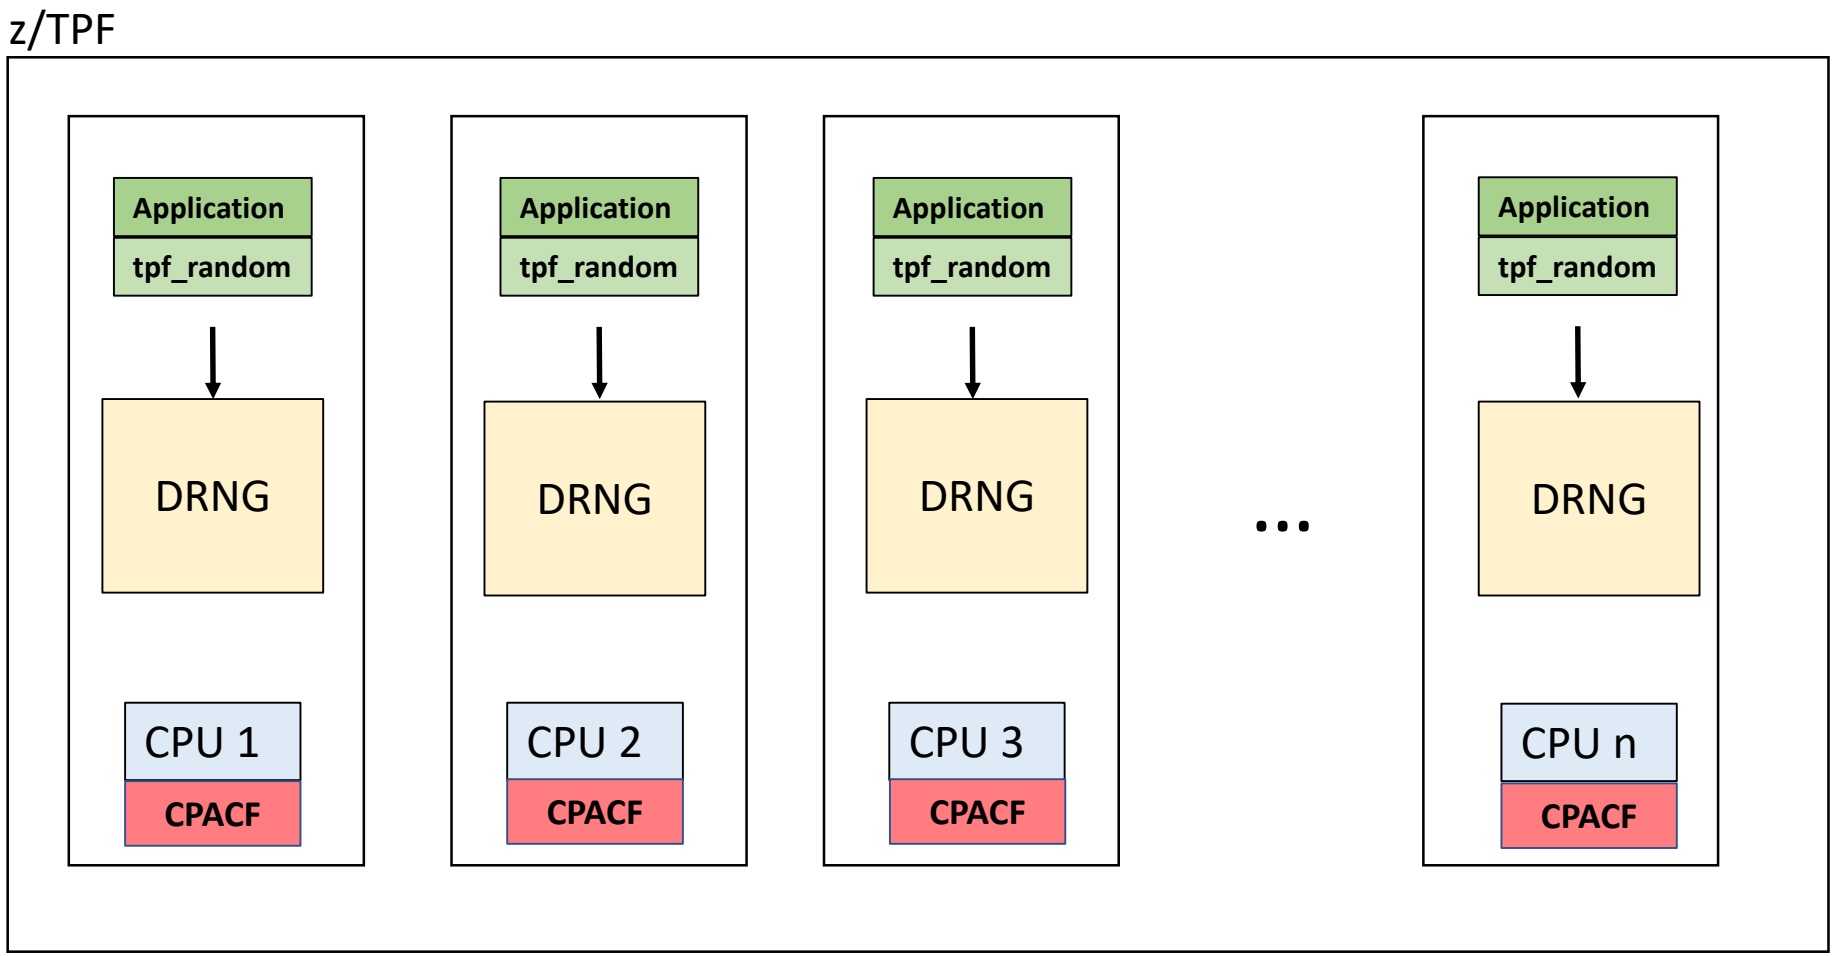

- For performance reasons, a separate DRNG exists on each CPU on the z/TPF processor.
- Reseeding occurs independently on each of the CPUs

Hardware Generated Random Numbers Performance

- **Tested on a z14 (700 series) – dedicated I-Streams**
- **Generating 64 bytes of random data**

- **Using the default RANDSEED value of 1024**
  - **Over 1 million 64-byte random numbers generated per second on each I-Stream**

Hardware Generated Random Numbers Summary

- Scheduled for delivery in 2Q 2019, PJ45130
- **Random numbers generated through the tpf_random function conform to the standards put in place by the National Institute of Standards and Technology (NIST) special publication 800-90A**

512 bit SHA-2 Message Digests

What is a Message Digest?

- Message digests are one-way cryptographic hash algorithms

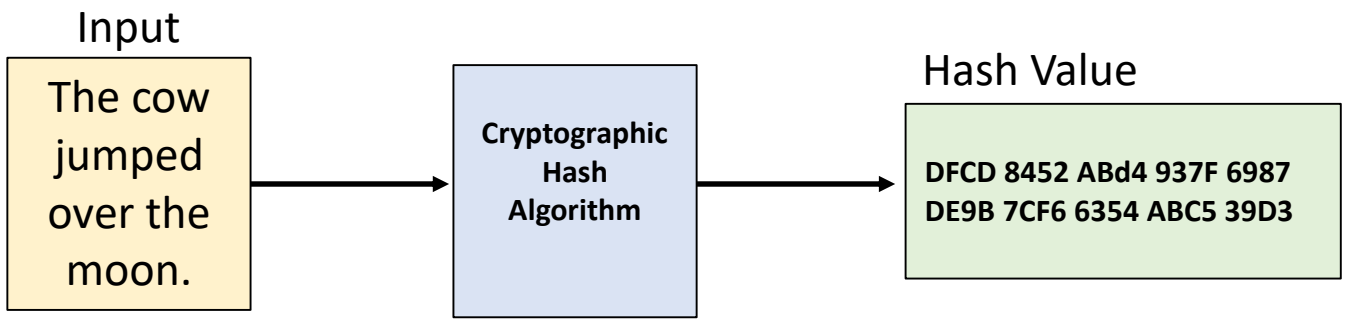

- Same message always produces the same hash
- Cannot generate the message given the hash value
- Small change to the message will result in completely different hash values

Current z/TPF Message Digest APIs

- Message digests on z/TPF are primarily done using the Secure Hash Algorithm – SHA
  - SHA-1, 160-bit (20 byte) hash created
    - Can use the tpf_SHA1_xxx APIs to generate SHA-1 hashes
      - For example, tpf_SHA1_Digest()
  - 256-bit SHA-2, 256-bit (32 byte) hash created
    - Can use the tpf_SHA2_xxx APIs to generate 256-bit SHA-2 hashes
      - For example, tpf_SHA256_Digest()
- The z/TPF message digest APIs require the necessary CPACF support to be enabled on the processor.

Current z/TPF Message Digest Support

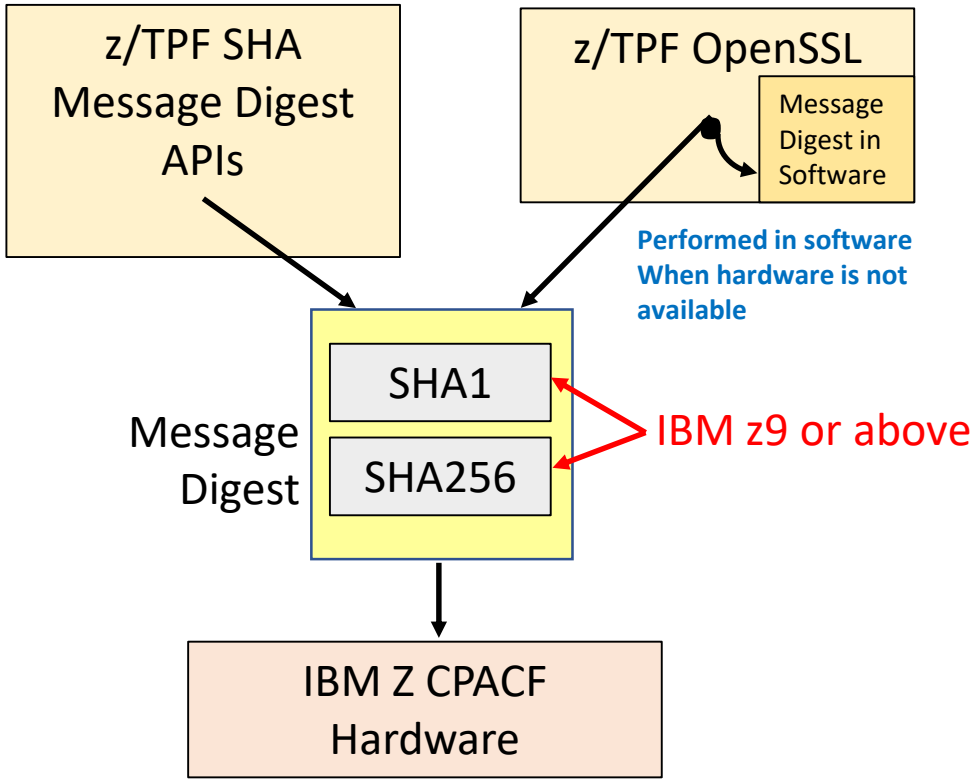

- Message digests are created via z/TPF API or through z/TPF OpenSSL
  - OpenSSL will perform message digests in software when hardware is unavailable.
- The SHA-1 and SHA-256 message digest algorithms available on IBM z9 or above.

Introducing SHA-512 to z/TPF

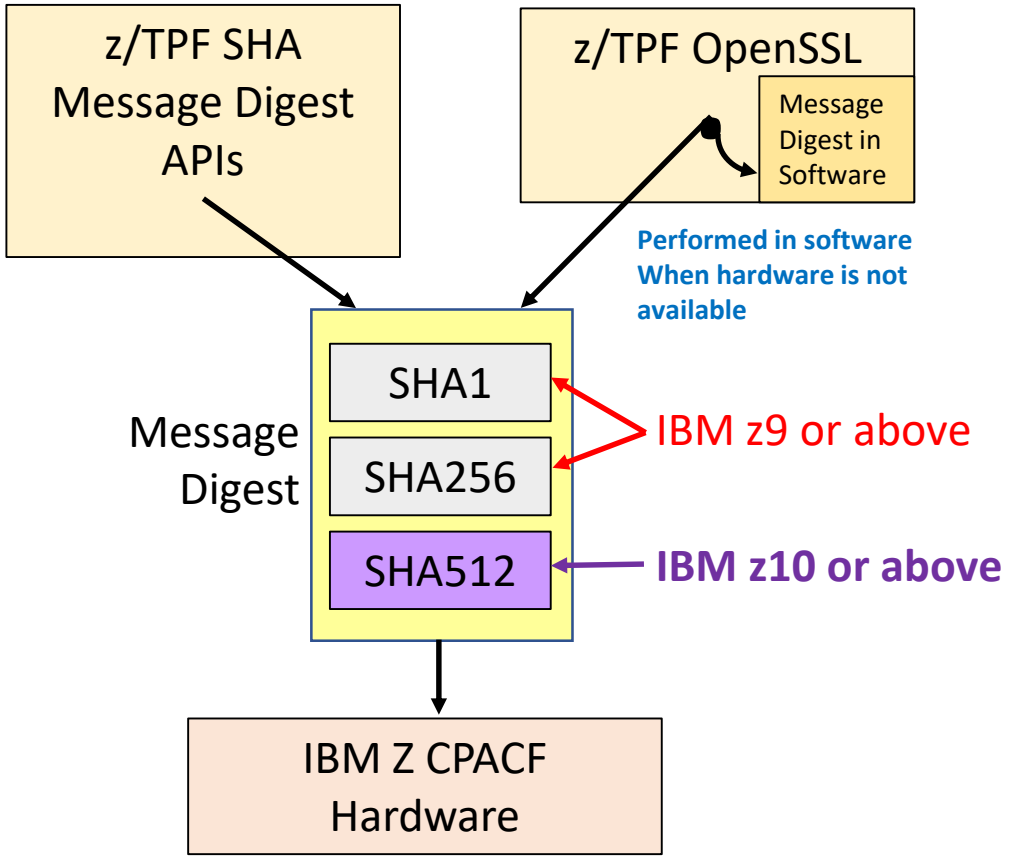

- New APIs to create 512-bit SHA-2 message digests.
- Hardware support is available on IBM z10 or higher.
- The OpenSSL community has not adopted the use of 512-bit SHA digests in its cipher suite….yet!

SHA-512 Support on z/TPF

- New APIs to create 512-bit message digests
  - tpf_SHA512_xxx

- Scheduled for delivery in 2Q 2019, PJ45130

CPACF Performance Enhancements

Use of CPACF Hardware is Growing Rapidly

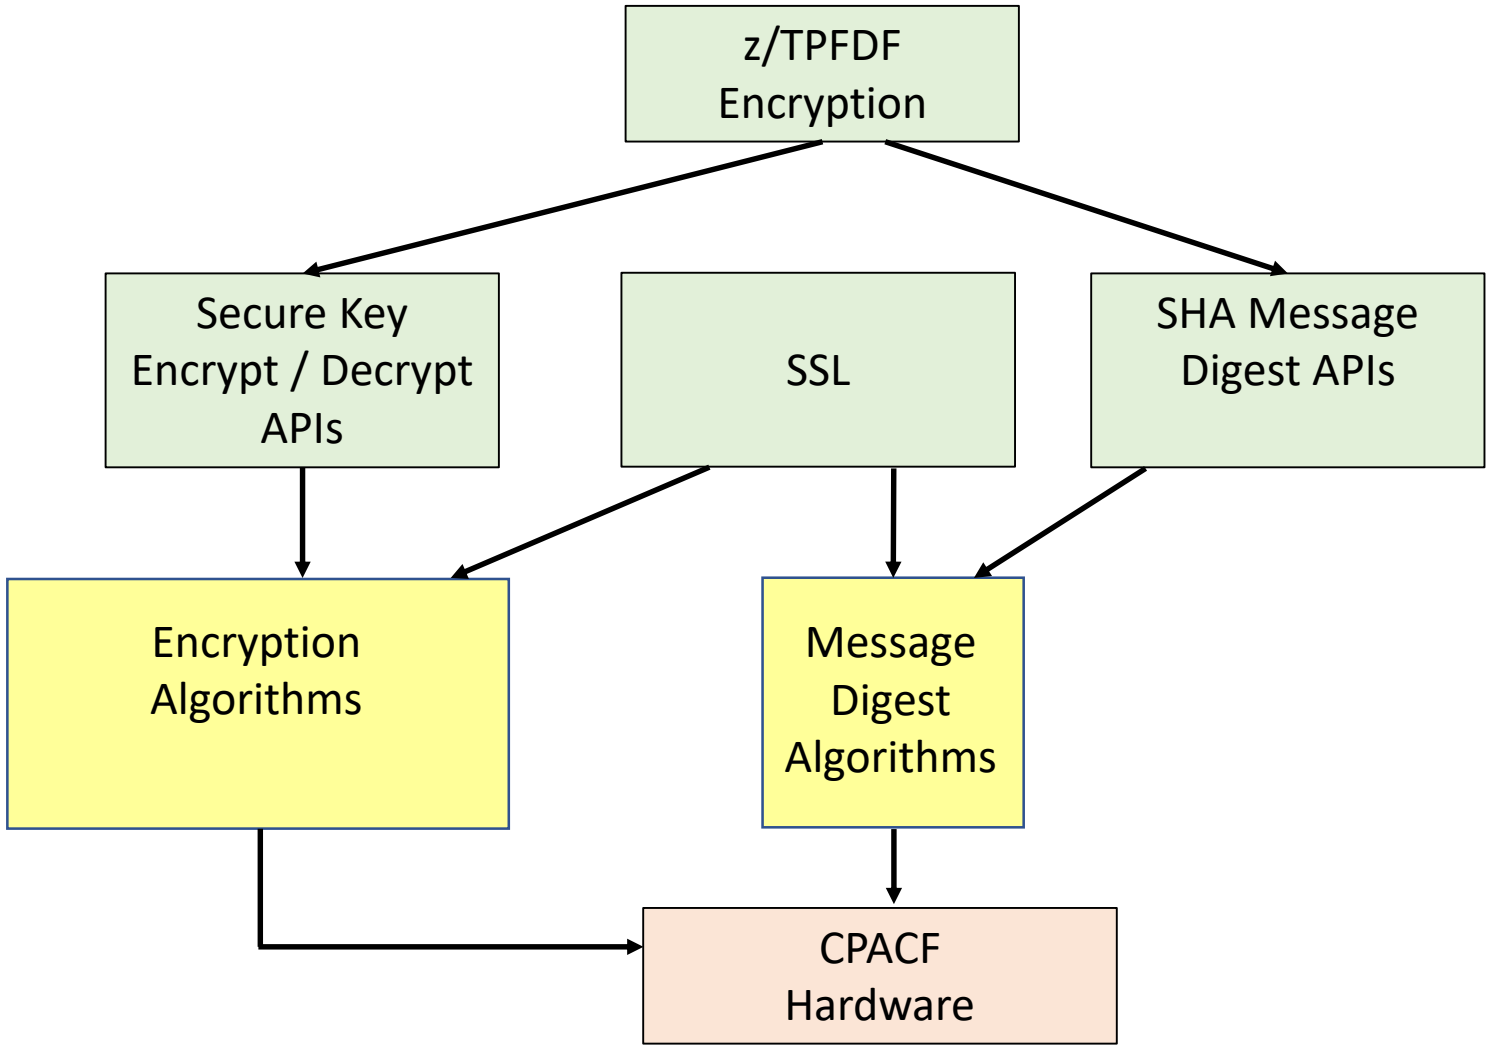

- Security and encryption is growing
- Critical that z/TPF software path to drive CPACF hardware requests is streamline for performance.

CPACF Performance Improvements

- The following performance improvements were made:
  - Reduced linkage cost through repackaging
  - Remove locking on the CPACF statistical table
  - Remove unnecessary ECB heap requests

- These performance changes would effect
  - z/TPF message digest APIs - tpf_SHAxxx
  - z/TPF secure key encryption APIs
    - tpf_encrypt_data/tpf_decrypt_data
  - OpenSSL operations performed in CPACF
  - z/TPFDF encryption

CPACF Performance Results

- **tpf_encrypt_data and tpf_decrypt_data APIs**
  - Up to a 17% CPU reduction encrypting/decrypting 4K of data using the AES-128 cipher algorithm.
  - Up to a 14% CPU reduction encrypting/decrypting 4K of data using the AES-256 cipher algorithm.
- **tpf_SHAxxx_digest APIs**
  - Up to a 10% CPU reduction creating SHA-256 digests on 4K of data.
- **OpenSSL Processing**
  - Up to 18% CPU reduction of time spent in OpenSSL to process a message

- Scheduled for delivery in 2Q 2019, PJ45130

Recap

- TCP/IP Performance Improvements – **PJ45093**, 3Q 2018
- Increasing the number of OSA read buffers – **PJ45555**, 1Q 2019
- z/TPF hardware generated random numbers – **PJ45130**, 2Q 2019
- 512-bit SHA-2 Message Digests – **PJ45130**, 2Q 2019
- CPACF Performance Improvements – **PJ45130**, 2Q 2019

Thank You!

Questions or 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.