DPDK

DATA PLANE DEVELOPMENT KIT

Flexible And Extensible Support for New Protocol Processing with DPDK using Dynamic Device Personalization

Brian Johnson, Intel® Corporation

Andrey Chilikin, Intel® Corporation

DPDK Summit - San Jose – 2017

Network Landscape

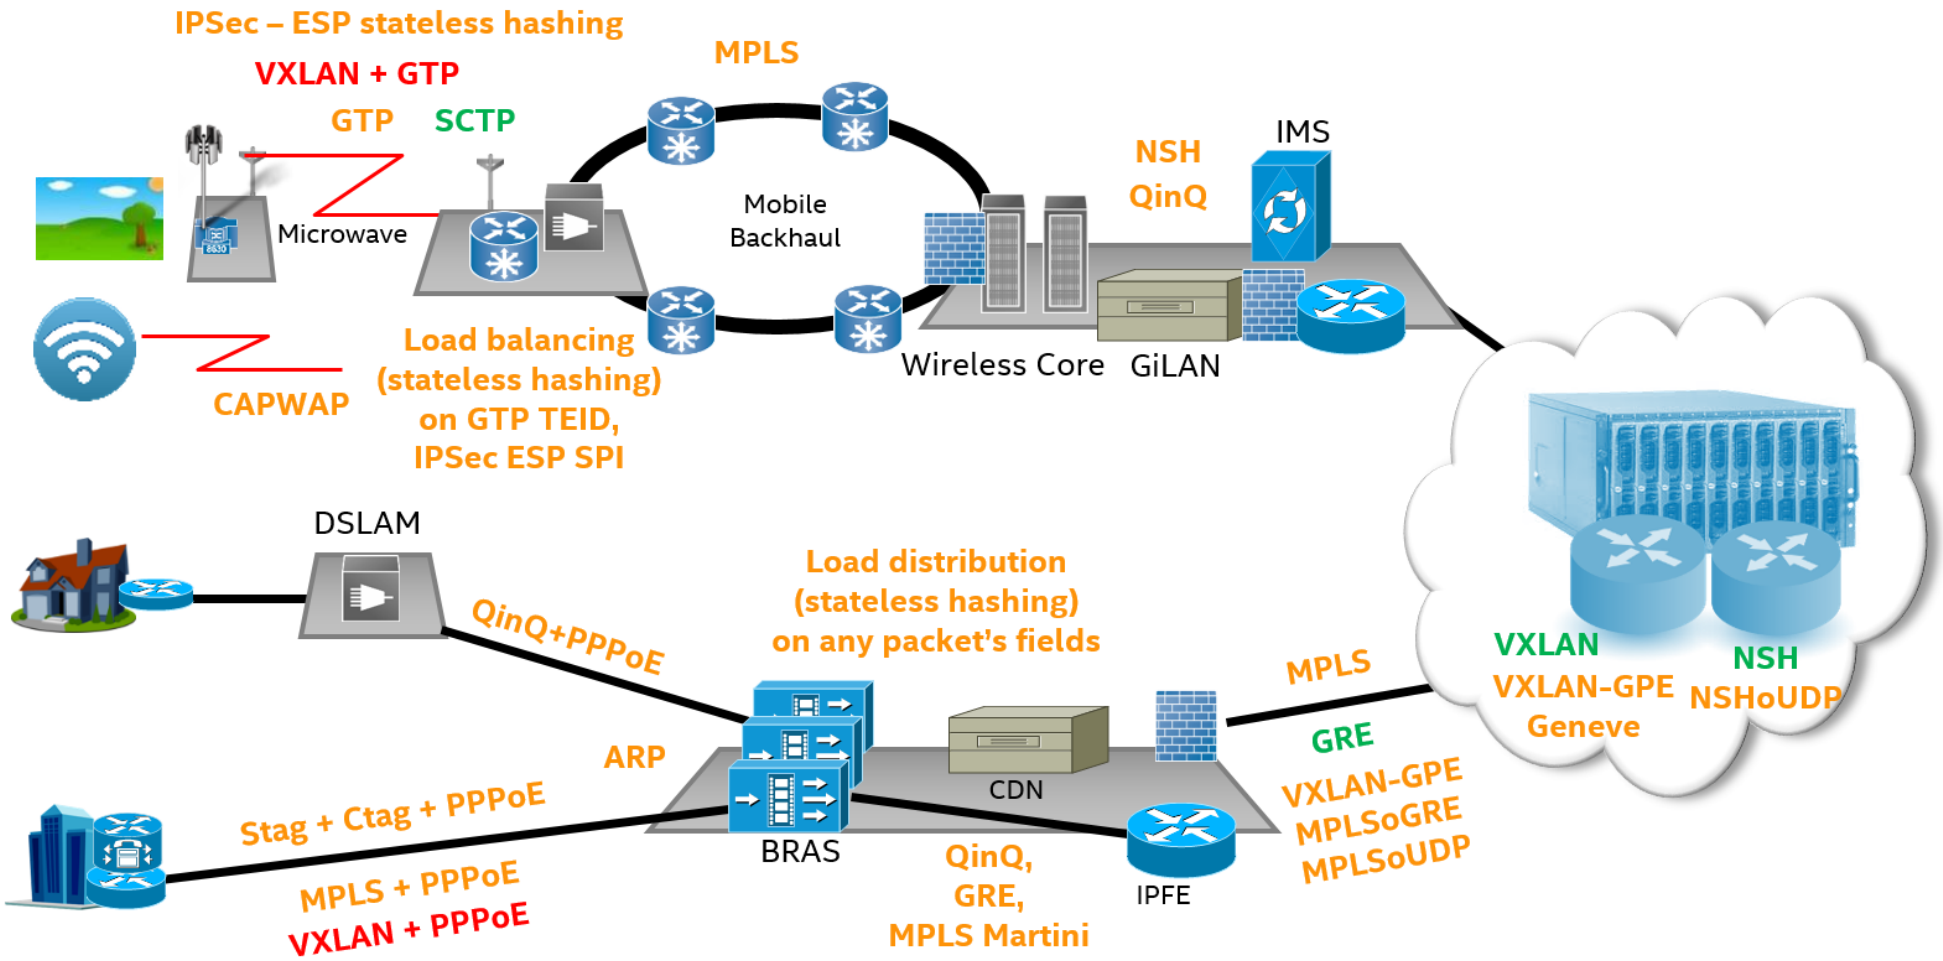

IPv4, IPv6, TCP, UDP, SCTP, VXLAN, NVGRE, GRE

Network Landscape

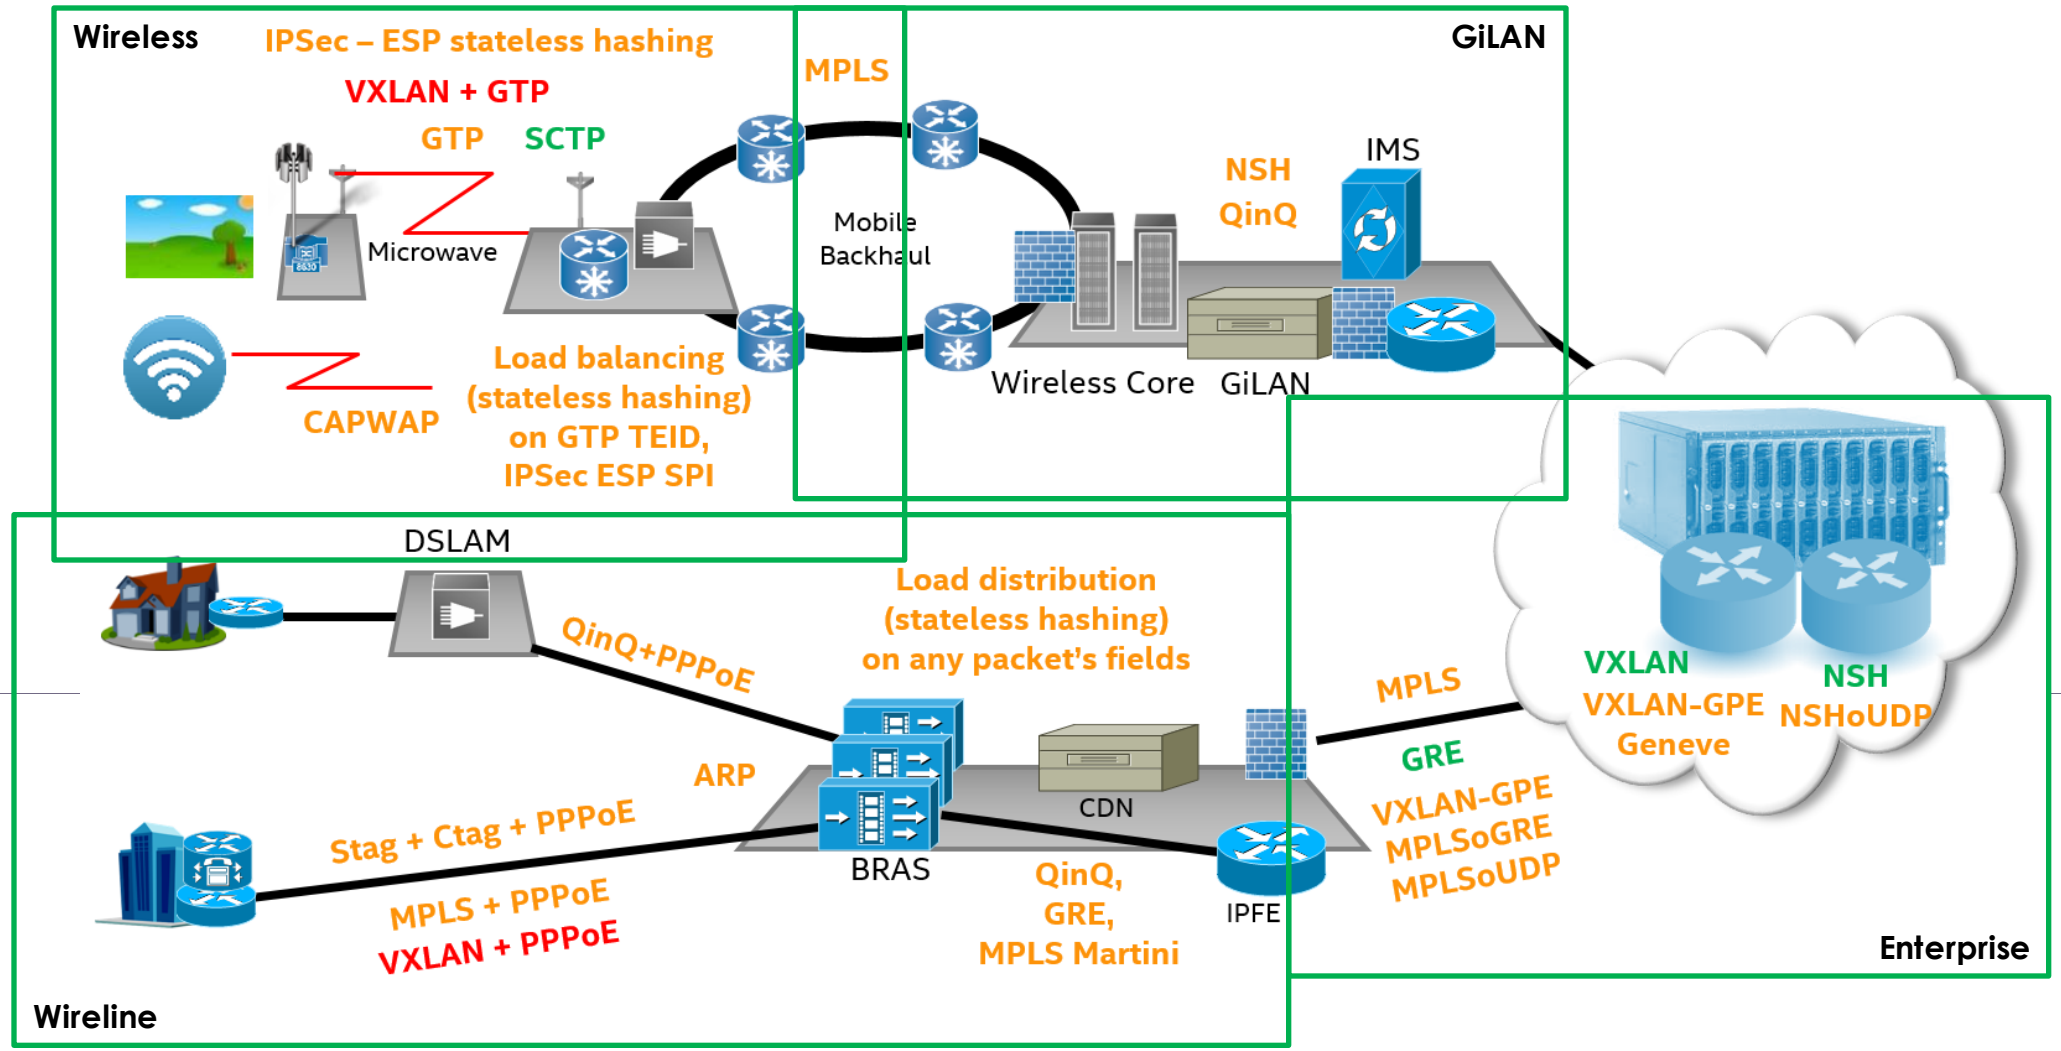

Dynamic Device Personalization (DDP)

Dynamic Device Personalization (DDP) profiles for Intel® Ethernet 700 Series enables run-time updating and configuration of parse graph to expand or modify protocol support allowing early adoption of new technologies

Example:

- ***Default Tunneling protocol support***
  - VXLAN, GENEVE, NVGRE, MPLS, VXLAN-GPE currently supported but limited by hardware resources
- ***Dynamic Device Personalization enables***
  1. Loadable profiles for packet classification for additional or modification of existing protocols
     - e.g. MPLSoGRE, GTP-U/GTP-C, PPPoE, QUIC, ESP, AH
  2. Configurable tunnel filters for traffic steering
     - Steering packets to VM queues on QinQ/tunnel ID

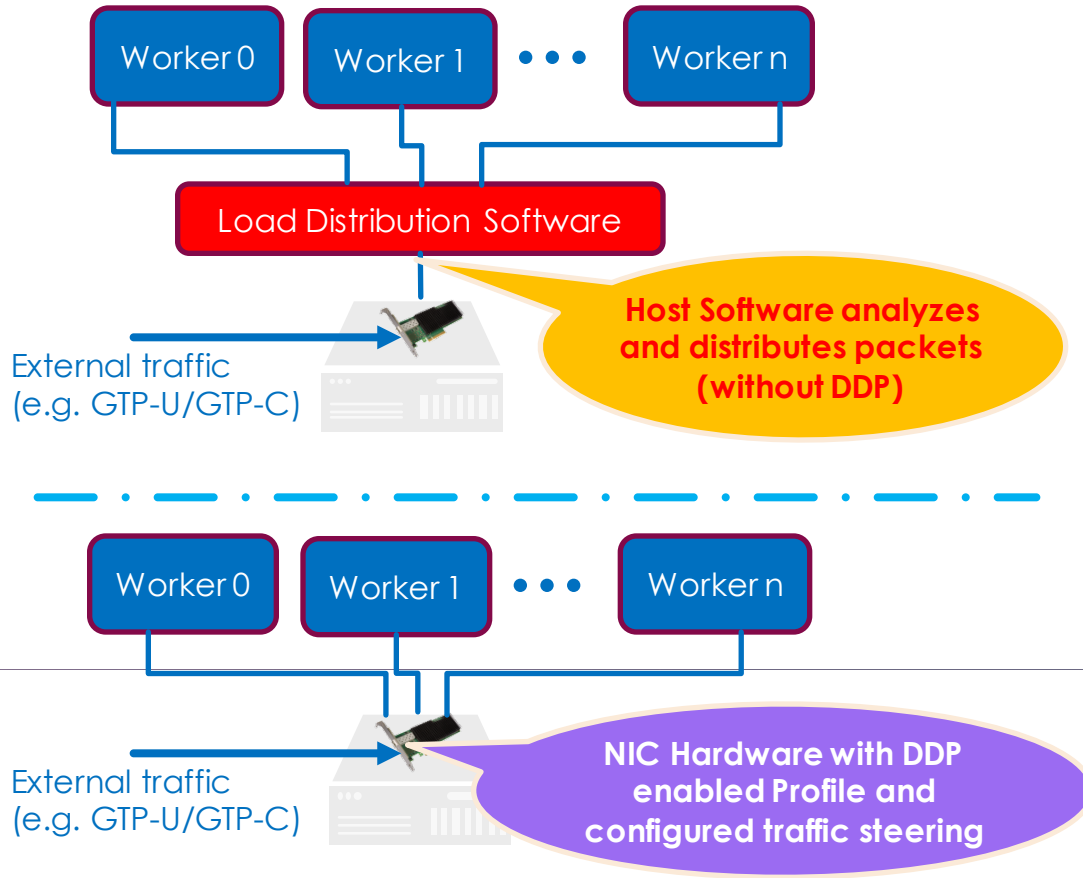

Packet analyzes and distribution is done on NIC with additional Rx Descriptor Data

Dynamic Device Personalization DPDK APIs*

- **rte_pmd_i40e_get_ddp_list()**
  get the list of registered profile(s)
- **rte_pmd_i40e_process_ddp_package()**
  download dynamic device personalization profile to the device and register profile
  rollback dynamic device personalization profile and un-register it
- **rte_pmd_i40e_get_ddp_info()**
  request information about a profile without downloading it to the NIC

* Requires XL710 firmware version 6.01
https://downloadcenter.intel.com/product/75021/Intel-Ethernet-Controller-XL710-Series

Dynamic Device Personalization: Without vs. With, GTP-U example

Default FVL6 processing (without Personalization profile):

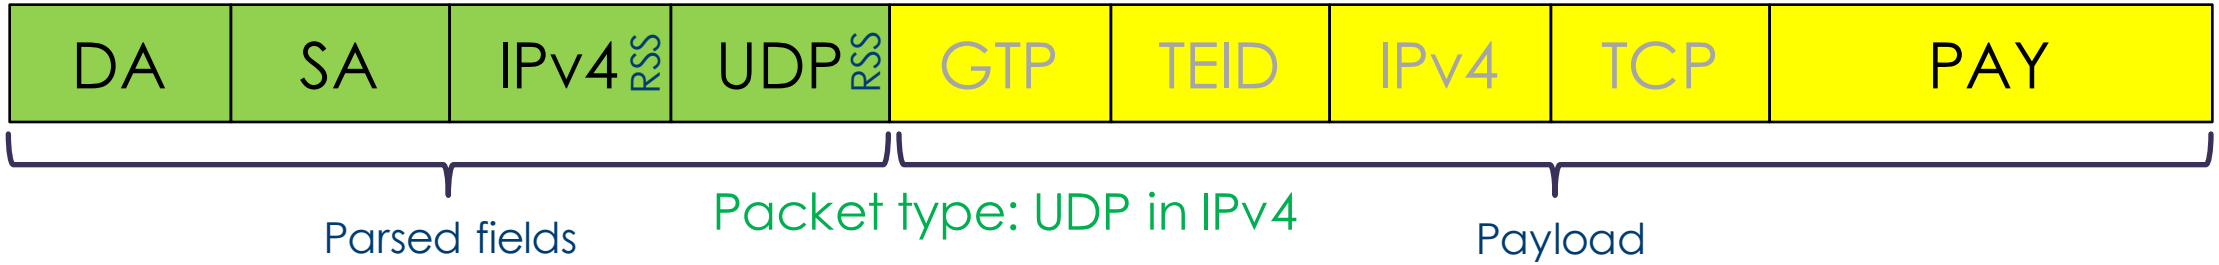

GTP-U is unknown flow type, so no RSS, FDIR or other filters are possible on encapsulated frame

FVL6 with Personalization profile applied:

GTP-U flow type is defined, encapsulated frame fields (including GTP TEID) can be used for RSS, FDIR. Encapsulated frame type is indicated on RX descriptor, for example, TCP in GTP-U or GTP-U echo message

GTP-C flow type is defined as well and has separate RSS/FDIR configuration.

About

DDP

Without Dynamic Device Personalization Profiles Loaded

DDP DEMO

DYNAMIC DEVICE PERSONALIZATION

- APPLY GTP PROFILE
- RSS PROFILE CONTROL

TRAFFIC CONTROLLS

- START ALL TRAFFIC
- STOP ALL TRAFFIC
- QUICS
- QUICL
- SCTP
- GTPU
- GTPC
- GTPU IPV4 UDP
- GTPU IPV4 TCP
- IPV4ESP

TRAFFIC AND TOOL STATUS

- QUICS
- QUICL
- SCTP
- GTPU
- GTPC
- GTPU IPV4 UDP
- GTPU IPV4 TCP
- IPv4ESP
- Traffic Generator
- Onboard Controller Port 1
- Onboard Controller Port 2
- Add-on Controller Port 1
- Add-on Controller Port 2

Onboard Controller Port 1

Q	PPS	PKT Type	Description
0	0	Reserved	NA
1	0	Reserved	NA
2	0	Reserved	NA
3	0	Reserved	NA
4	0	Reserved	NA
5	0	Reserved	NA
6	0	Reserved	NA
7	0	Reserved	NA
8	0	Reserved	NA
9	0	Reserved	NA
10	1734221	UDP	IPV4 UDP PAY4
11	0	Reserved	NA
12	0	Reserved	NA
13	0	Reserved	NA
14	0	Reserved	NA
15	0	Reserved	NA
Total:	1734907	RSS Applied	No

Onboard Controller Port 2

Q	PPS	PKT Type	Description
0	203316	UDP	IPV4 UDP PAY4
1	9054	SCTP	IPV4 SCTP PAY4
2	13412	SCTP	IPV4 SCTP PAY4
3	13787	SCTP	IPV4 SCTP PAY4
4	13794	SCTP	IPV4 SCTP PAY4
5	13786	SCTP	IPV4 SCTP PAY4
6	13827	SCTP	IPV4 SCTP PAY4
7	13777	SCTP	IPV4 SCTP PAY4
8	11603	SCTP	IPV4 SCTP PAY4
9	8974	SCTP	IPV4 SCTP PAY4
10	202881	UDP	IPV4 UDP PAY4
11	9228	SCTP	IPV4 SCTP PAY4
12	13619	SCTP	IPV4 SCTP PAY4
13	13453	SCTP	IPV4 SCTP PAY4
14	13646	SCTP	IPV4 SCTP PAY4
15	13738	SCTP	IPV4 SCTP PAY4
Total:	581889	RSS Applied	No

Add-on Controller Port 1

Q	PPS	PKT Type	Description
0	94053	UDP	IPV4 UDP PAY4
1	96614	UDP	IPV4 UDP PAY4
2	95479	UDP	IPV4 UDP PAY4
3	91824	UDP	IPV4 UDP PAY4
4	95923	UDP	IPV4 UDP PAY4
5	94799	UDP	IPV4 UDP PAY4
6	95579	UDP	IPV4 UDP PAY4
7	92408	UDP	IPV4 UDP PAY4
8	1604682	UDP	IPV4 UDP PAY4
9	1606796	UDP	IPV4 UDP PAY4
10	94826	UDP	IPV4 UDP PAY4
11	95692	UDP	IPV4 UDP PAY4
12	93381	UDP	IPV4 UDP PAY4
13	93275	UDP	IPV4 UDP PAY4
14	92738	UDP	IPV4 UDP PAY4
15	94679	UDP	IPV4 UDP PAY4
Total:	4533103	RSS Applied	No

Add-on Controller Port 2

Q	PPS	PKT Type	Description
0	0	Reserved	NA
1	0	Reserved	NA
2	0	Reserved	NA
3	0	Reserved	NA
4	0	Reserved	NA
5	0	Reserved	NA
6	0	Reserved	NA
7	1755741	PAY3	IPV4 PAY3
8	0	Reserved	NA
9	0	Reserved	NA
10	0	Reserved	NA
11	0	Reserved	NA
12	0	Reserved	NA
13	0	Reserved	NA
14	0	Reserved	NA
15	0	Reserved	NA
Total:	1756169	RSS Applied	No

Onboard Controller Profile

0 profile registered

Add-on Controller Profile

0 profile registered

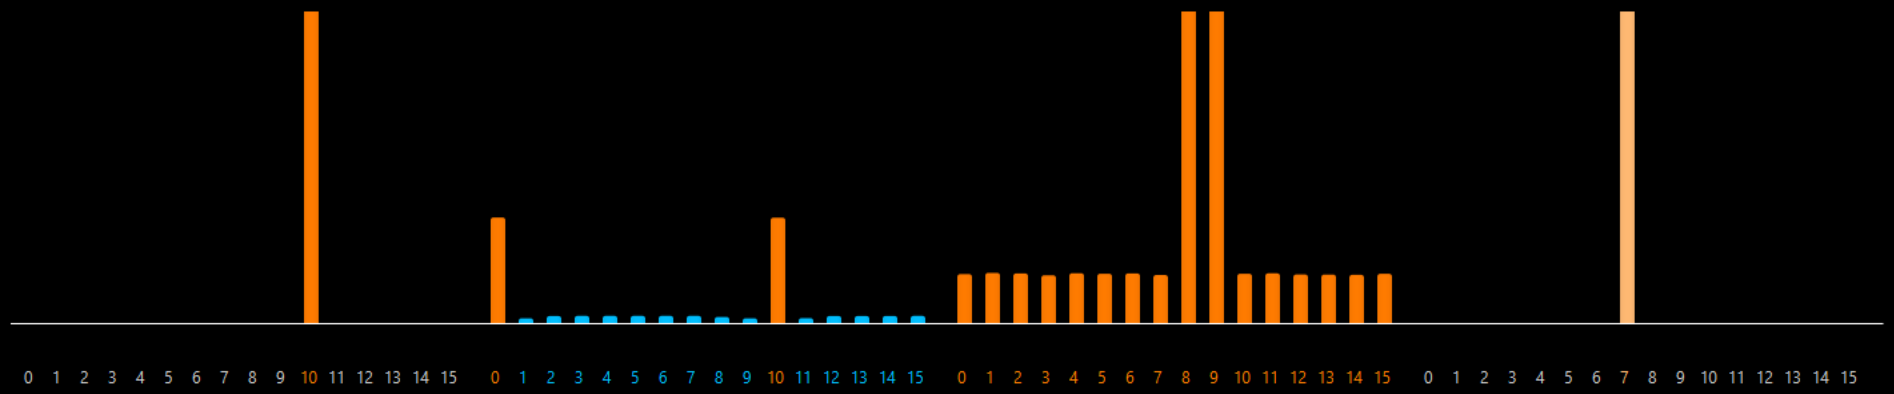

About

DDP

With Dynamic Device Personalization GTP, QUIC & IPsec Profiles Loaded

DDP DEMO

DYNAMIC DEVICE PERSONALIZATION

- ROLLBACK GTP PROFILE
- RSS PROFILE CONTROL

TRAFFIC CONTROLLS

- START ALL TRAFFIC
- STOP ALL TRAFFIC
- QUICS
- QUICL
- SCTP
- GTPU
- GTPC
- GTPU IPV4 UDP
- GTPU IPV4 TCP
- IPV4ESP

TRAFFIC AND TOOL STATUS

- QUICS
- QUICL
- SCTP
- GTPU
- GTPC
- GTPU IPV4 UDP
- GTPU IPV4 TCP
- IPv4ESP
- Traffic Generator
- Onboard Controller Port 1
- Onboard Controller Port 2
- Add-on Controller Port 1
- Add-on Controller Port 2

Onboard Controller Port 1

Q	PPS	PKT Type	Description
0	109803	GTPU	IPV4 GTPU IPV4 UDP PAY4
1	111005	GTPU	IPV4 GTPU IPV4 UDP PAY4
2	110373	GTPU	IPV4 GTPU IPV4 UDP PAY4
3	111493	GTPU	IPV4 GTPU IPV4 UDP PAY4
4	111977	GTPU	IPV4 GTPU IPV4 TCP PAY4
5	112500	GTPU	IPV4 GTPU IPV4 TCP PAY4
6	109899	GTPU	IPV4 GTPU IPV4 TCP PAY4
7	109680	GTPU	IPV4 GTPU IPV4 UDP PAY4
8	106058	GTPU	IPV4 GTPU IPV4 TCP PAY4
9	110066	GTPU	IPV4 GTPU IPV4 TCP PAY4
10	108017	GTPU	IPV4 GTPU IPV4 UDP PAY4
11	113080	GTPU	IPV4 GTPU IPV4 TCP PAY4
12	111865	GTPU	IPV4 GTPU IPV4 UDP PAY4
13	108941	GTPU	IPV4 GTPU IPV4 UDP PAY4
14	110746	GTPU	IPV4 GTPU IPV4 TCP PAY4
15	107464	GTPU	IPV4 GTPU IPV4 UDP PAY4
Total:	1763030	RSS Applied	No

Onboard Controller Port 2

Q	PPS	PKT Type	Description
0	33264	GTPC	IPV4 GTPC PAY4
1	33335	GTPC	IPV4 GTPC PAY4
2	37331	SCTP	IPV4 SCTP PAY4
3	37783	GTPC	IPV4 GTPC PAY4
4	37551	SCTP	IPV4 SCTP PAY4
5	37858	GTPU	IPV4 GTPU PAY4
6	38314	GTPU	IPV4 GTPU PAY4
7	38206	GTPC	IPV4 GTPC PAY4
8	36148	GTPC	IPV4 GTPC PAY4
9	33721	GTPU	IPV4 GTPU PAY4
10	33588	GTPU	IPV4 GTPU PAY4
11	33554	GTPU	IPV4 GTPU PAY4
12	37606	GTPU	IPV4 GTPU PAY4
13	38011	GTPC	IPV4 GTPC PAY4
14	37794	GTPU	IPV4 GTPU PAY4
15	37845	SCTP	IPV4 SCTP PAY4
Total:	582007	RSS Applied	No

Add-on Controller Port 1

Q	PPS	PKT Type	Description
0	189427	QUICS	IPV4 UDP QUICS PAY4
1	185592	QUICS	IPV4 UDP QUICS PAY4
2	689886	QUICNOCID	IPV4 UDP QUICNOCID PAY4
3	186839	QUICS	IPV4 UDP QUICS PAY4
4	189628	QUICL	IPV4 UDP QUICL PAY4
5	188432	QUICL	IPV4 UDP QUICL PAY4
6	186860	QUICL	IPV4 UDP QUICL PAY4
7	190149	QUICL	IPV4 UDP QUICL PAY4
8	190520	QUICS	IPV4 UDP QUICS PAY4
9	697372	QUICS	IPV4 UDP QUICS PAY4
10	184565	QUICS	IPV4 UDP QUICS PAY4
11	695082	QUICS	IPV4 UDP QUICS PAY4
12	189895	QUICS	IPV4 UDP QUICS PAY4
13	189719	QUICS	IPV4 UDP QUICS PAY4
14	192575	QUICS	IPV4 UDP QUICS PAY4
15	186045	QUICL	IPV4 UDP QUICL PAY4
Total:	4532131	RSS Applied	No

Add-on Controller Port 2

Q	PPS	PKT Type	Description
0	109728	IPv4ESP	IPv4 ESP PAY4
1	109748	IPv4ESP	IPv4 ESP PAY4
2	109748	IPv4ESP	IPv4 ESP PAY4
3	109728	IPv4ESP	IPv4 ESP PAY4
4	109734	IPv4ESP	IPv4 ESP PAY4
5	109734	IPv4ESP	IPv4 ESP PAY4
6	109730	IPv4AH	IPv4 AH PAY4
7	109729	IPv4AH	IPv4 AH PAY4
8	109728	IPv4ESP	IPv4 ESP PAY4
9	109747	IPv4ESP	IPv4 ESP PAY4
10	109728	IPv4ESP	IPv4 ESP PAY4
11	109748	IPv4ESP	IPv4 ESP PAY4
12	109744	IPv4ESP	IPv4 ESP PAY4
13	109734	IPv4ESP	IPv4 ESP PAY4
14	109728	IPv4AH	IPv4 AH PAY4
15	109728	IPv4AH	IPv4 AH PAY4
Total:	1755677	RSS Applied	No

Onboard Controller Profile

1 profile registered

track	version	ow name
80000008	0.0.0.9	00 GTP-C/GTP-U IPv4/IPv6 payload

Add-on Controller Profile

1 profile registered

track	version	ow name
DE000003	0.0.0.3	00 XL710 QUIC/IPSec Demo

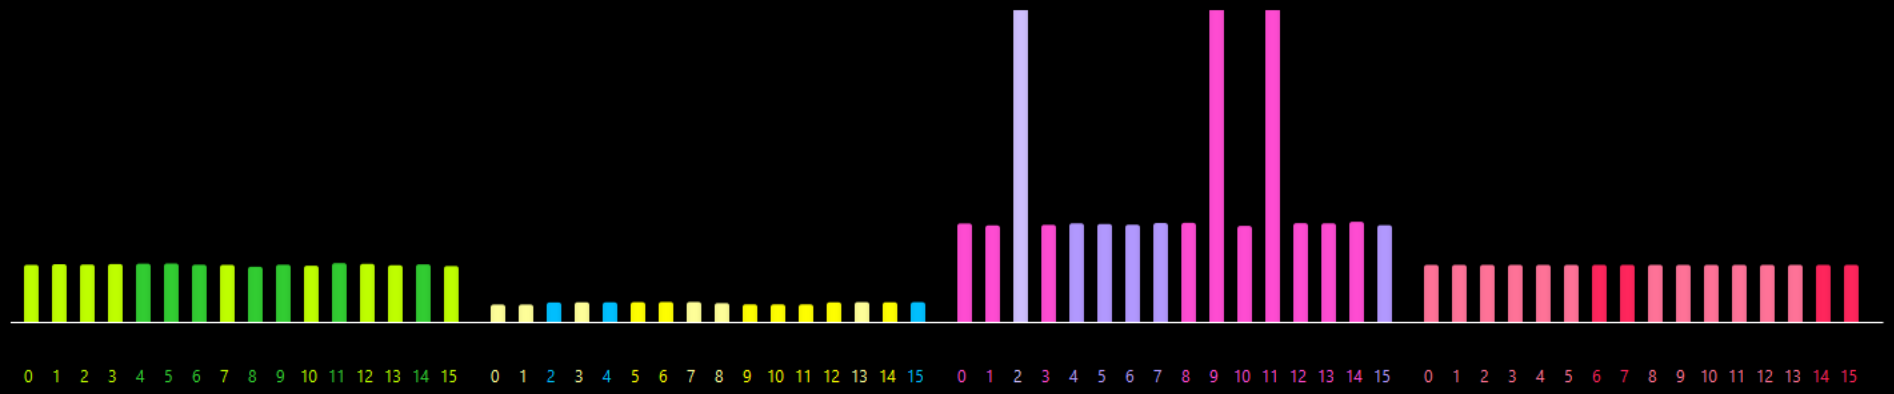

About

DDP

With Dynamic Device Personalization RSS Queue Regions Assigned

DDP DEMO

DYNAMIC DEVICE PERSONALIZATION

- ROLLBACK GTP PROFILE
- RSS PROFILE CONTROL

TRAFFIC CONTROLLS

- START ALL TRAFFIC
- STOP ALL TRAFFIC
- QUICS
- QUICL
- SCTP
- GTPU
- GTPC
- GTPU IPV4 UDP
- GTPU IPV4 TCP
- IPV4ESP

TRAFFIC AND TOOL STATUS

- QUICS
- QUICL
- SCTP
- GTPU
- GTPC
- GTPU IPV4 UDP
- GTPU IPV4 TCP
- IPv4ESP
- Traffic Generator
- Onboard Controller Port 1
- Onboard Controller Port 2
- Add-on Controller Port 1
- Add-on Controller Port 2

Onboard Controller Port 1

Q	PPS	PKT Type	Description
0	109754	GTPU	IPV4 GTPU IPV4 UDP PAY4
1	110952	GTPU	IPV4 GTPU IPV4 TCP PAY4
2	110402	GTPU	IPV4 GTPU IPV4 TCP PAY4
3	111378	GTPU	IPV4 GTPU IPV4 UDP PAY4
4	112025	GTPU	IPV4 GTPU IPV4 TCP PAY4
5	112493	GTPU	IPV4 GTPU IPV4 TCP PAY4
6	109908	GTPU	IPV4 GTPU IPV4 UDP PAY4
7	109699	GTPU	IPV4 GTPU IPV4 UDP PAY4
8	106098	GTPU	IPV4 GTPU IPV4 UDP PAY4
9	110149	GTPU	IPV4 GTPU IPV4 UDP PAY4
10	107978	GTPU	IPV4 GTPU IPV4 UDP PAY4
11	113027	GTPU	IPV4 GTPU IPV4 TCP PAY4
12	111885	GTPU	IPV4 GTPU IPV4 UDP PAY4
13	108932	GTPU	IPV4 GTPU IPV4 TCP PAY4
14	110747	GTPU	IPV4 GTPU IPV4 TCP PAY4
15	107333	GTPU	IPV4 GTPU IPV4 UDP PAY4
Total:	1762017	RSS Applied	No

Onboard Controller Port 2

Q	PPS	PKT Type	Description
0	0	Reserved	NA
1	0	Reserved	NA
2	98235	SCTP	IPV4 SCTP PAY4
3	95826	SCTP	IPV4 SCTP PAY4
4	96437	GTPU	IPV4 GTPU PAY4
5	97407	GTPU	IPV4 GTPU PAY4
6	96923	GTPC	IPV4 GTPC PAY4
7	97150	GTPC	IPV4 GTPC PAY4
8	0	Reserved	NA
9	0	Reserved	NA
10	0	Reserved	NA
11	0	Reserved	NA
12	0	Reserved	NA
13	0	Reserved	NA
14	0	Reserved	NA
15	0	Reserved	NA
Total:	581954	RSS Applied	Yes

Add-on Controller Port 1

Q	PPS	PKT Type	Description
0	0	Reserved	NA
1	0	Reserved	NA
2	503584	QUICNOCID	IPV4 UDP QUICNOCID PAY4
3	503606	QUICNOCID	IPV4 UDP QUICNOCID PAY4
4	503603	QUICNOCID	IPV4 UDP QUICNOCID PAY4
5	0	Reserved	NA
6	379318	QUICL	IPV4 UDP QUICL PAY4
7	380911	QUICL	IPV4 UDP QUICL PAY4
8	373296	QUICL	IPV4 UDP QUICL PAY4
9	377442	QUICL	IPV4 UDP QUICL PAY4
10	0	Reserved	NA
11	0	Reserved	NA
12	380172	QUICS	IPV4 UDP QUICS PAY4
13	376661	QUICS	IPV4 UDP QUICS PAY4
14	377023	QUICS	IPV4 UDP QUICS PAY4
15	376995	QUICS	IPV4 UDP QUICS PAY4
Total:	4532841	RSS Applied	Yes

Add-on Controller Port 2

Q	PPS	PKT Type	Description
0	109764	IPv4AH	IPv4 AH PAY4
1	109764	IPv4AH	IPv4 AH PAY4
2	109762	IPv4AH	IPv4 AH PAY4
3	109762	IPv4AH	IPv4 AH PAY4
4	109784	IPv4AH	IPv4 AH PAY4
5	109762	IPv4AH	IPv4 AH PAY4
6	109784	IPv4AH	IPv4 AH PAY4
7	109784	IPv4AH	IPv4 AH PAY4
8	109786	IPv4ESP	IPv4 ESP PAY4
9	109764	IPv4ESP	IPv4 ESP PAY4
10	109784	IPv4AH	IPv4 AH PAY4
11	109762	IPv4AH	IPv4 AH PAY4
12	109762	IPv4ESP	IPv4 ESP PAY4
13	109783	IPv4AH	IPv4 AH PAY4
14	109783	IPv4AH	IPv4 AH PAY4
15	109763	IPv4AH	IPv4 AH PAY4
Total:	1756188	RSS Applied	No

Onboard Controller Profile

1 profile registered

track	version	ow name
80000008	0.0.0.9	00 GTP-C/GTP-U IPv4/IPv6 payload

Add-on Controller Profile

1 profile registered

track	version	ow name
DE000003	0.0.0.3	00 XL710 QUIC/IPSec Demo

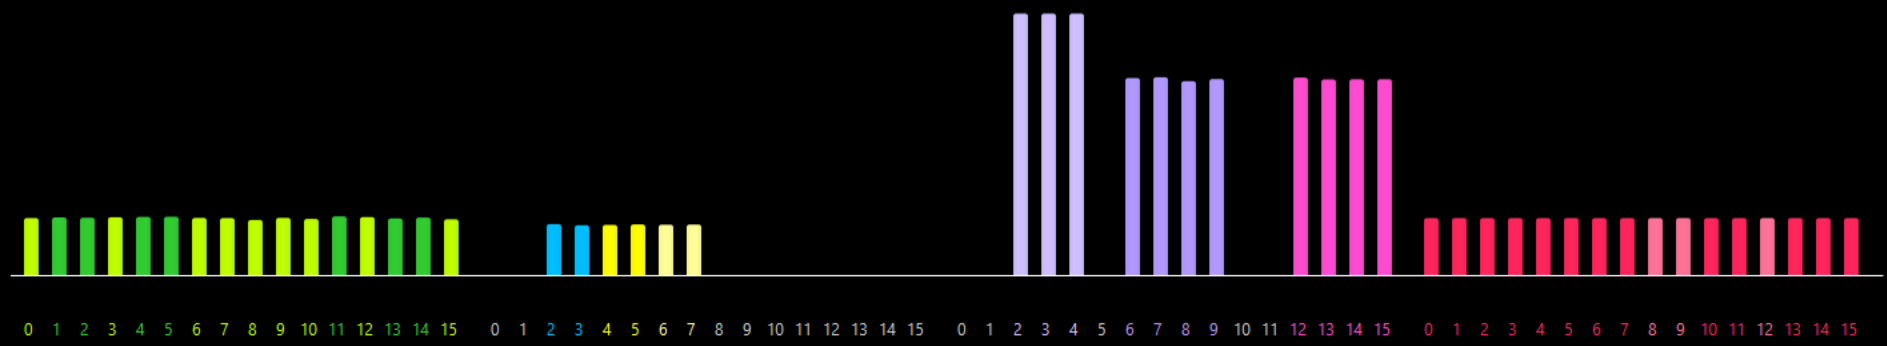

Summary

- Dynamic Device Personalization can be used on any OS supported by DPDK
  - For example, QUIC protocol classification on both Windows and Linux OSes

DPDK on Microsoft Windows

DPDK on Linux

- Dynamic Device Personalization allows a DPDK application to enable identification of new protocols, for example, GTP, PPPoE, QUIC, MPLSoGRE, SRv6, etc without requiring a reboot due to changing hardware or firmware updates.
- Allows for existing Parse Graphs to be modified to support new usage models as they are defined enabling faster time to market for DPDK enable applications.

Thank You

Dynamic Device Personalization

Core Team

Eyal Prashker
Yochai Hagvi
Andrey Chilikin
Brian Johnson
Robin Giller

DPDK Team

Helin Zhang
Heqing Zhu

Dynamic Device Personalization with GTP and QUIC Profiles

Special Thanks to *Michael Shearer* and *Patrick Kutch* for all the help on the development of the DDP Dashboard and server instrumentation development
BIF Framework -- https://github.com/intel/Board-Instrumentation-Framework