This presentation covers Gen-Z Control Space and Control structures.

Disclaimer

This document is provided 'as is' with no warranties whatsoever, including any warranty of merchantability, noninfringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. Gen-Z Consortium disclaims all liability for infringement of proprietary rights, relating to use of information in this document. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Gen-Z is a trademark or registered trademark of the Gen-Z Consortium.

All other product names are trademarks, registered trademarks, or servicemarks of their respective owners.

All material is subject to change at any time at the discretion of the Gen-Z Consortium

http://genzconsortium.org/

Control Space Basics

- Control Space can contain up to 2^{52} bytes of addressable resources
  - Data Space can contain up to 2^{64} bytes of addressable resources
- Control Space is distinct from Data Space
  - Partitioning prevents erroneous or malicious access to critical resources and to simplify management
- At a minimum, Control Space contains control structures used for configuration and management
- Control Space may also contain:
  - Component-specific structures and resources
  - Accelerator or embedded processing executables
  - Translation and mapping tables
  - Working set space
  - Etc.
- Control Space can be accessed using an out-of-band 2-wire interconnect and / or using in-band Control OpClass packets
- Control Space pages may be mapped into a Requester ZMMU to enable memory-speed access and control

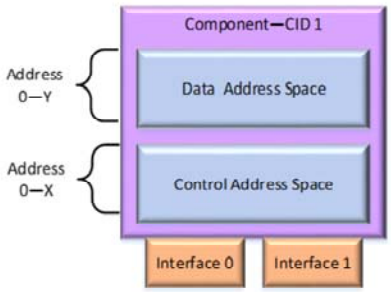

Gen-Z architecture supports two address space—the Data Address Space and the Control Address Space. The Data Address Space is used by applications and contains application-specific data. It supports up to 2^{64} bytes of addressable resources. Control Space is used by Gen-Z configuration and management and supports up to 2^{52} bytes of addressable resources.

Control Space is independently addressed to prevent erroneous or malicious access to critical resources. Further, this simplifies management and security solutions.

At a minimum, Control Space contains control structures used by Gen-Z configuration and management. It can also contain component-specific structures and resources, embedded executables, translation and mapping tables used by memory and storage media services, etc.

Control Space can be accessed using an out-of-band 2-wire interconnect. For example, to simplify memory solutions, P2P-Core memory uses only out-of-band management interconnect. Similarly, to dramatically simplify 802.3 electrical implementations, out-of-band management is used. If components support the Control OpClass, then components use in-band management.

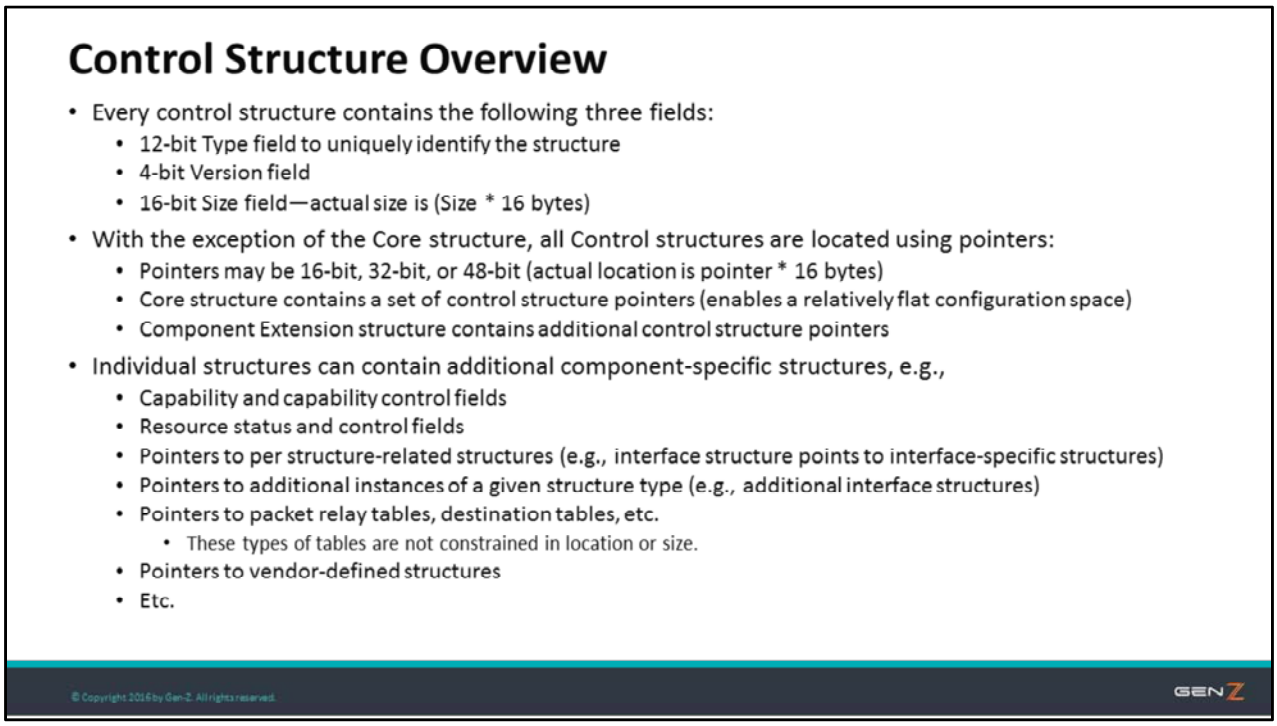

Control structures are used to configure and manage a Gen-Z component. Every structure contains three standardized fields: Type, Version, and Size.

The Core structure is located at Control Space address byte 0. The Core structure is used to locate all other control structures. It contains a set of pointer fields. There are three pointer sizes: 16b, 32b, and 48b. A pointer is used to locate byte 0 of a control structure. The actual location is (pointer * 16B). If an implementation requires additional pointers, then the Component Extension structure is used; it is a structure of pointers. Multiple Component Extension structures can be provisioned.

Individual structures contain a variety of component-specific structures. Some structures contain pointers to structure-related structures, e.g., the Interface structure contains pointers to all interface-specific structures. Some structures contain pointers to additional instances of a given structure, e.g., if a component supports multiple interfaces, then the additional interfaces are located using a Next Interface Pointer (Next I-PTR). Structures can also contain pointers to non-control structures, e.g., packet relay tables, destination tables, etc. These tables will vary in size, and can be larger than the maximum size of a control structure.

Mandatory Control Structures

- The following structures are mandatory in every component:
  - Core
    - 512B structure located at Control Space byte 0
    - Structure contains:
      - Common configuration and management fields
      - Pointers fields used to access other control structures
      - UUIDs used to identify and manage the component
  - OpCode Set
    - Used to identify supported OpClasses and OpCodes
    - Used to enable OpClasses and OpCodes to ensure interoperability between communicating components
    - May support multiple OpCode Set structures to support different component needs or protocol generations
  - Interface
    - One Interface structure per component interface
    - Used to manage the interface and to access all interface-specific structures
  - Interface PHY
    - One Interface PHY structure per supported physical layer per component interface
    - Used to configure and manage supported physical layer(s)
- Mandatory structures require less than 1 KiB of Control Space

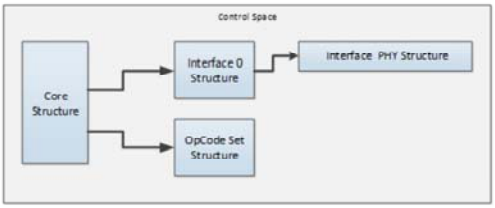

GEN Z

Every component supports at least the 4 mandatory structures—the Core, OpCode Set, Interface, and Interface PHY structure.

The Core structure is the base or root structure used to access all other structures. It contains the most common configuration and management fields, pointers to access other control structures, and UUIDs used to identify and manage the component.

The OpCode Set structure is used to communicate the supported OpClasses and OpCodes. Using this information, management enables the common OpClasses and OpCodes in communicating components to ensure interoperability.

There is one Interface structure per component interface. It is used to configure and manage the interface and to access all Interface-specific structures.

The Interface PHY structure is used to manage the physical layer associated with the interface. If an interface supports multiple physical layers, then there is on Interface PHY structure per physical layer.

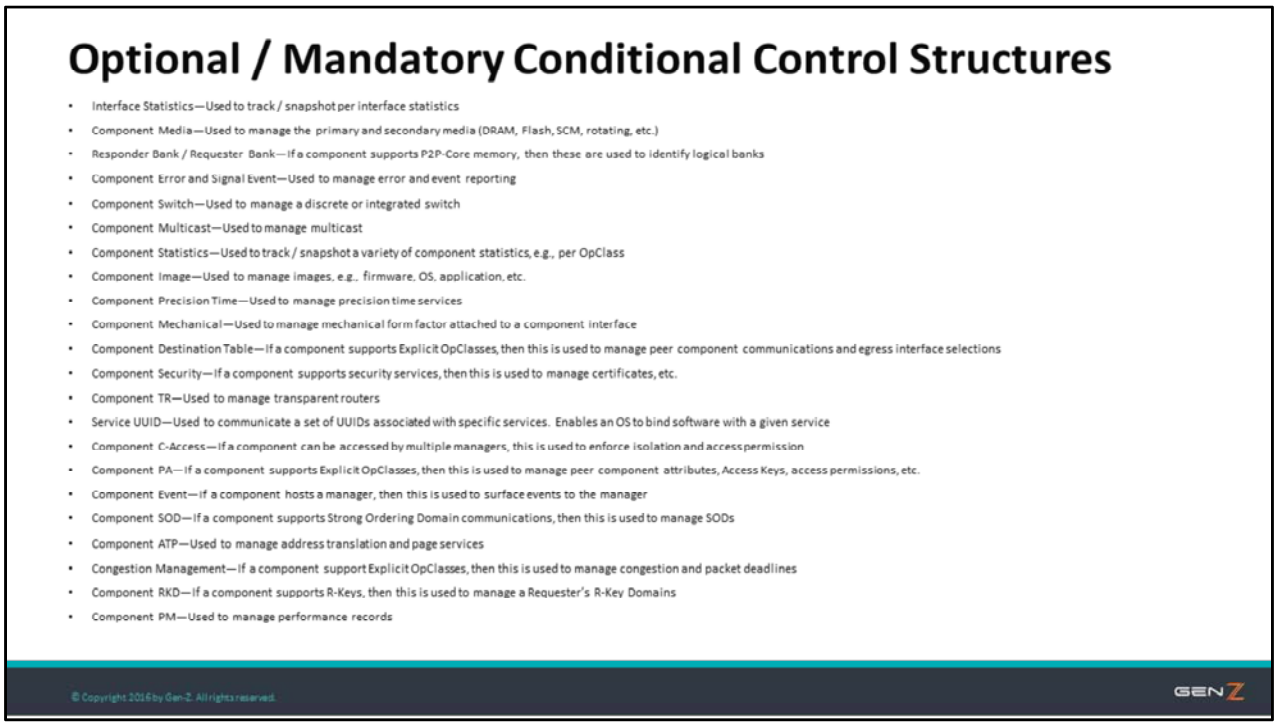

Gen-Z supports a set of optional and mandatory conditional structures. Optional structures are implemented at the component's discretion. Mandatory conditional structures are mandatory if specific conditions are met. For example, if a component support the P2P-Core OpClass, then based on its role, it supports the Responder Bank structure or the Requester Bank structure. These structures are used to manage P2P-Core memory logical banks.

Control Structure Organization

- With the exception of the Core structure, control structures can be placed anywhere in Control Space
- If a component is managed by the Primary Manager, e.g., point-to-point memory component, then control structures can be packed to minimize Control Space resources.
- If a component can be managed by multiple managers, e.g., Primary Manager / Fabric Manager, an OS, a resource manager, etc., then control structures should be grouped to enable access control / permission protection. For example, the following groups are specified:
  - Baseline group—Core, OpCode Set, Interface, Interface PHY
  - Routing / Fabric group—structures used to configure packet relay tables and operations
  - Media group—structures used to manage addressable media
  - PCI Group—structures used to manage PCI / PCIe configuration headers
  - Vendor-defined—component-specific structures
  - Dedicated page—structures that may be owned by a resource manager, an OS, etc.

GEN Z

If a component is managed only by the Primary Manager, e.g., system firmware, then the control structures can be organized to minimize Control Space resources.

If a component can be managed by multiple managers, then the control structures should be grouped to enable fine-grain access control and permission. Gen-Z specifies multiple control structure groupings to simplify implementation and to ensure different managers can be granted read-only or read-write access or no access to a given set of control structures.

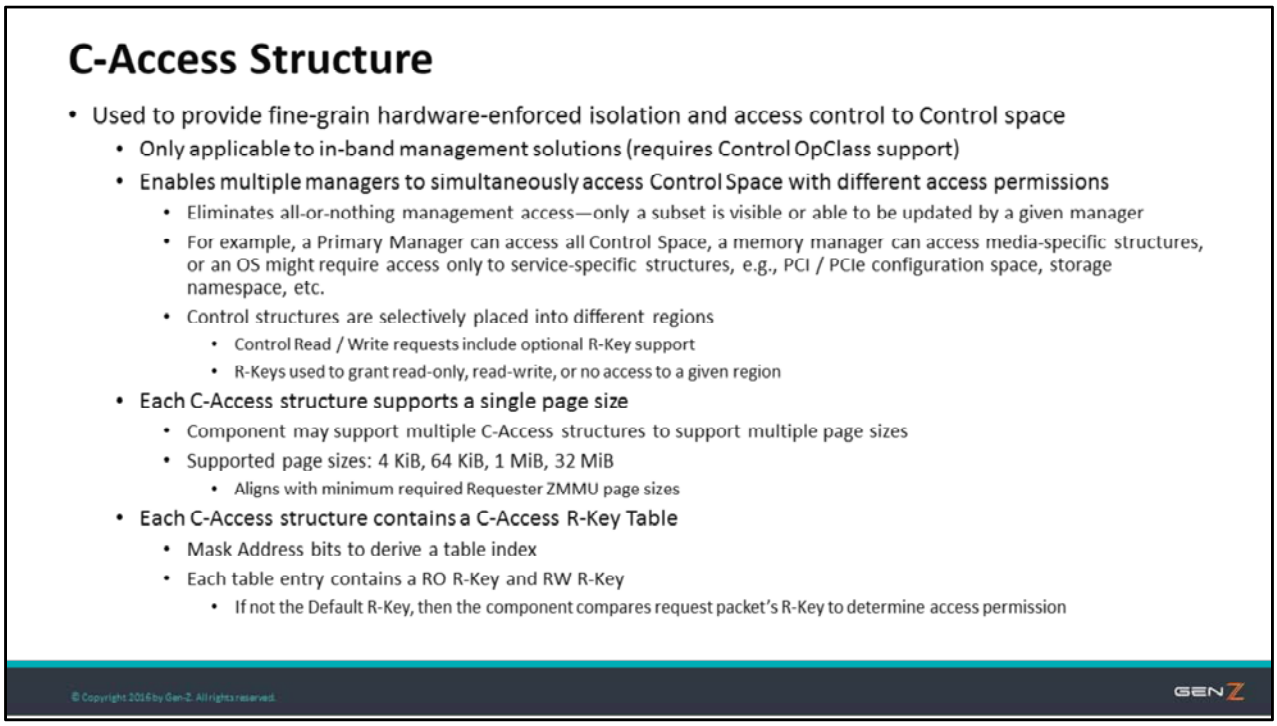

If a component can be managed by multiple managers, then it should support the Component C-Access structure. This structure should be part of the Baseline group, and configured by either the Primary Manager or the Primary or Secondary Fabric Manager.

A component provisions one C-Access structure per supported page size. Using a fixed page size per table simplifies hardware access to the C-Access R-Key Table—simply mask off a set of bits to use as an index into the table. Each table entry contains two R-Keys—the read-only and the read-write R-Keys. If these fields are not the Default R-Key, then the component performs R-Key validation to determine if access is permitted or not. Each manager can be provided different access permission.

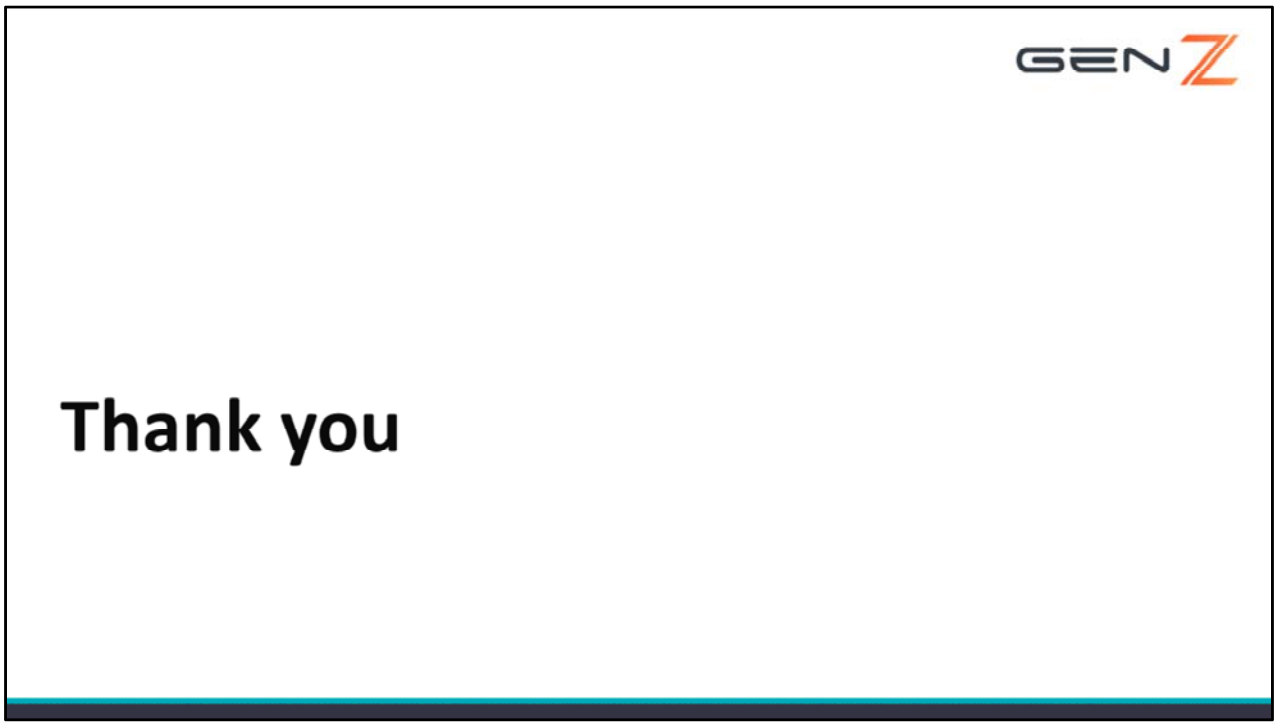

This concludes this presentation. Thank you.