

# Unleashing AI at Scale: Exploring the IBM Spyre Accelerator on IBM z17

Mitchell Bride  
mitch.bride@ibm.com  
IBM WSC Hardware Technical Specialist

# Contents

What is the IBM Spyre Accelerator?

Hardware Overview

Use Cases

Configuration

- Requirements
- Preconfiguration
- Installing ACC
- Installing SSA

SOD


Questions

# What is the IBM Spyre Accelerator?



# IBM Telum II Processor

- 5nm technology, 5.5GHz
- 8 cores with 20% area reduction and improved microprocessor power management
- 40% more cache per core
- **NEW:** On-chip Data Processing Unit (DPU): Increased I/O performance with 70% reduction in power for I/O management, RAS, reduced latency
- 2<sup>nd</sup>-gen AI Accelerator for high-speed inferencing with fine tuning
- 8x dedicated AI processing per core



# Dedicated low-latency Integrated Accelerator for AI


Centralized on-chip accelerator shared by all cores



With IBM z17, process up to 5 million inference operations per second with less than 1ms response time using a Credit Card Fraud Detection model\*



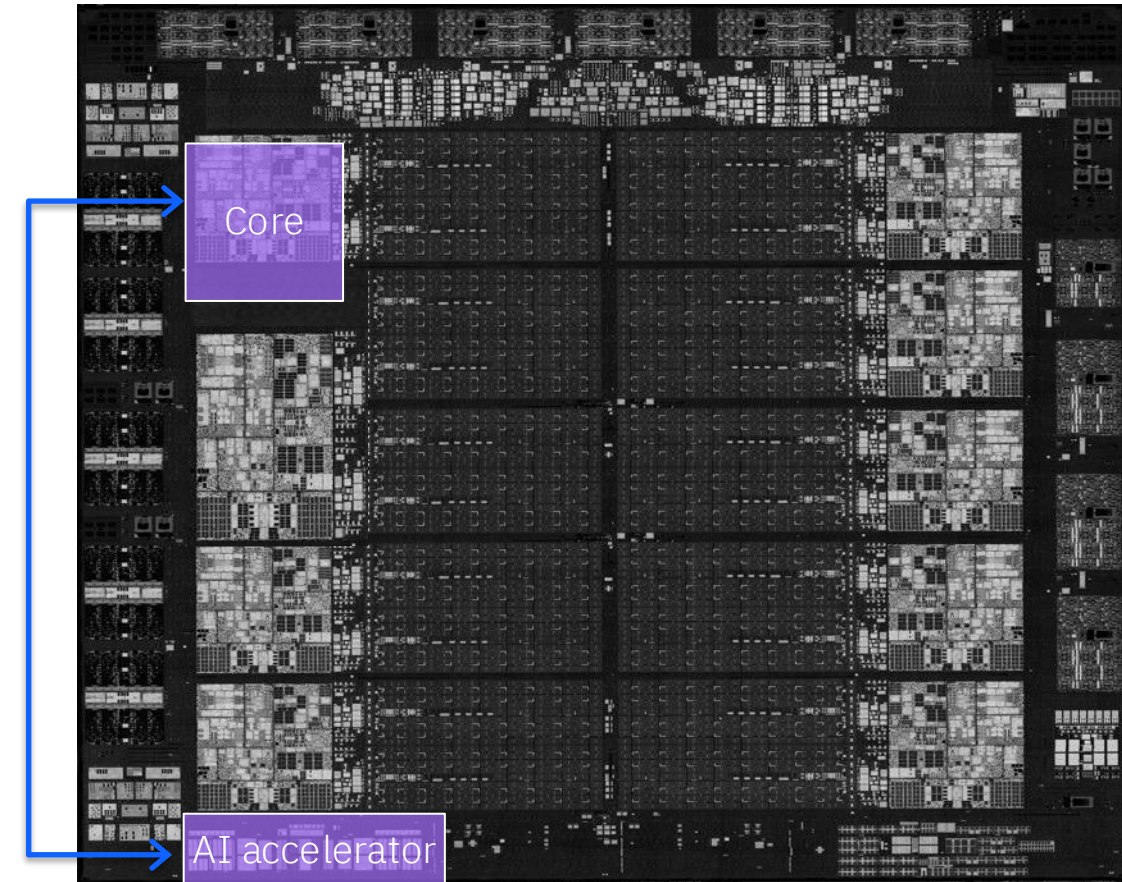
Very low and consistent inference latency and scalable capacity



New Function with the z17

- Transform function between different data layouts/types
- Gaussian Error Linear Unit function
- NNPA instruction can be dispatched to any available accelerator in a drawer

Every CP chip has one Integrated Accelerator for AI built-in



# More AI acceleration on IBM z17

## In-transaction AI with encoder LLMs and multiple AI model techniques

### 2nd Gen on-chip AI accelerator in Telum II


- Support for LLM compute primitives
- Improved quantization and matrix operations
- Improved AI processing over IBM z16<sup>9</sup>



## AI workload balancing during peak usage

### In-drawer intelligent routing

- Remote AI processing
- Up to 8x AI processing available



## Optimize generative AI and LLM use cases


### IBM Spyre Accelerator cards\*

- 32 Gen AI-ready cores per adapter card
- Up to 48 adapter cards per system



\*Available starting in 4Q25

# Predictive AI and generative AI on IBM z17 deliver unique capabilities



### Advantage

- Real-time AI inference
- In-transaction AI at scale
- Inference every transaction<sup>1</sup>

### Advantage

- Gen AI with IBM Z RAS
- Security and data protection
- Energy efficiency


## Multiple-model AI at speed and scale

Combine machine learning and LLMs to achieve higher prediction accuracy and expanded insights




# Unlock greater AI accuracy, productivity, and agility

- Apply multiple AI models to transactions for better accuracy with optimal latency
- Increase IT productivity
- Predict and reduce IT incidents for greater resiliency
- On-board new Z skills in half the time



Inferencing



Assistants



Agents

*Effective combination of AI inferencing, assistants\*, and agents\* to transform and simplify the mainframe experience*

\* Upon IBM Spyre Accelerator availability

AI on IBM z17  
makes more  
valuable  
outcomes  
possible  
for every  
industry

### Financial Services



FRAUD  
DETECTION



ANTI-MONEY  
LAUNDERING



RISK SCORING




CREDIT DECISIONING


### Insurance



CLAIMS FRAUD  
DETECTION



INTELLIGENT  
UNDERWRITING



TEXT/IMAGE  
PROCESSING



PRODUCT  
RECOMMENDATIONS


### Government




GEOSPATIAL IMAGE  
ANALYSIS



FRAUD & SENTIMENT  
ANALYSIS



AUDIT &  
COMPLIANCE




CHAT SERVICE


### Others



RETAIL  
INVENTORY/DEMAND  
FORECASTING



SYSTEM ADMIN  
ASSISTANT




CODE  
ASSISTANT




TRANSPORTATION  
LOGISTICS

# IBM watsonx Assistant delivers exceptional customer experiences



Accessible and scalable AI that you can rely on with your customers

Accurate and consistent customer experiences out of the box with no-code conversational AI features and LLM-powered algorithms that achieve higher accuracy with less training, driving down TCO, and speeding up time-to-value.




Assistants that can tell, show, and do

Every problem has an optimal solution. Guide your customers down the right path, including contextual, summarized answers and give your team the ability to deliver the best experience for each unique situation.




Deliver consistent and personalized experiences with flexible extensibility

Deliver your best support experience across all channels. Integrate LLM-powered experiences with the systems and processes that run your business, compatible with any tech stack.



Orient your team around true customer outcomes

Unify your analytics tools across the entire customer support journey. Empower your teams to experiment, analyze, and optimize your end-to-end customer experience.



Build and iterate quickly with no code and LLM-powered authoring

An easy-to-use build experience, pre-designed repeatable conversational patterns, and generative authoring accelerate builds without the need for IT or development resources.

# AI on IBM Z and IBM® LinuxONE

## Future Generative AI use cases



### Tax Document Processing

Leverage GenAI to summarize submitted tax documents and highlight key findings such as mismatched income reporting, unusually high deductions, etc.

### Financial Documents Processing

Leverage GenAI to summarize financial documents and business reports to extract key data points such as financial metrics, and performance indicators, and identify essential information needed for compliance processes like financial audits

### Information Search and Data Extraction

Extract key points from content that has been stored as unstructured text. Search and identify essential information needed for compliance processes like financial audits and regulatory compliance

### Customs Screening for Suspicious Cargo


Utilize Generative AI to identify potentially suspicious cargo through sophisticated image processing techniques and analysis of textual descriptions associated with each shipment

# IBM Spyre Accelerator PCIe attached card

- 75W PCIe gen5 x16 adapter
- 128GB of LPDDR5 memory
- 300+ TOPS
- Up to 8 cards per I/O drawer
- Generative AI: 8 cards form a logical cluster
- 1TB of aggregate memory
- 1.6TB per second aggregate memory bandwidth



Designed to handle **Large Language Models**  
and **Generative AI** use cases




# z17 PCIe attached AI accelerator placements

Frame C  
3 IO drawer

Frame B  
3 IO drawer

Frame Z  
5 IO drawer

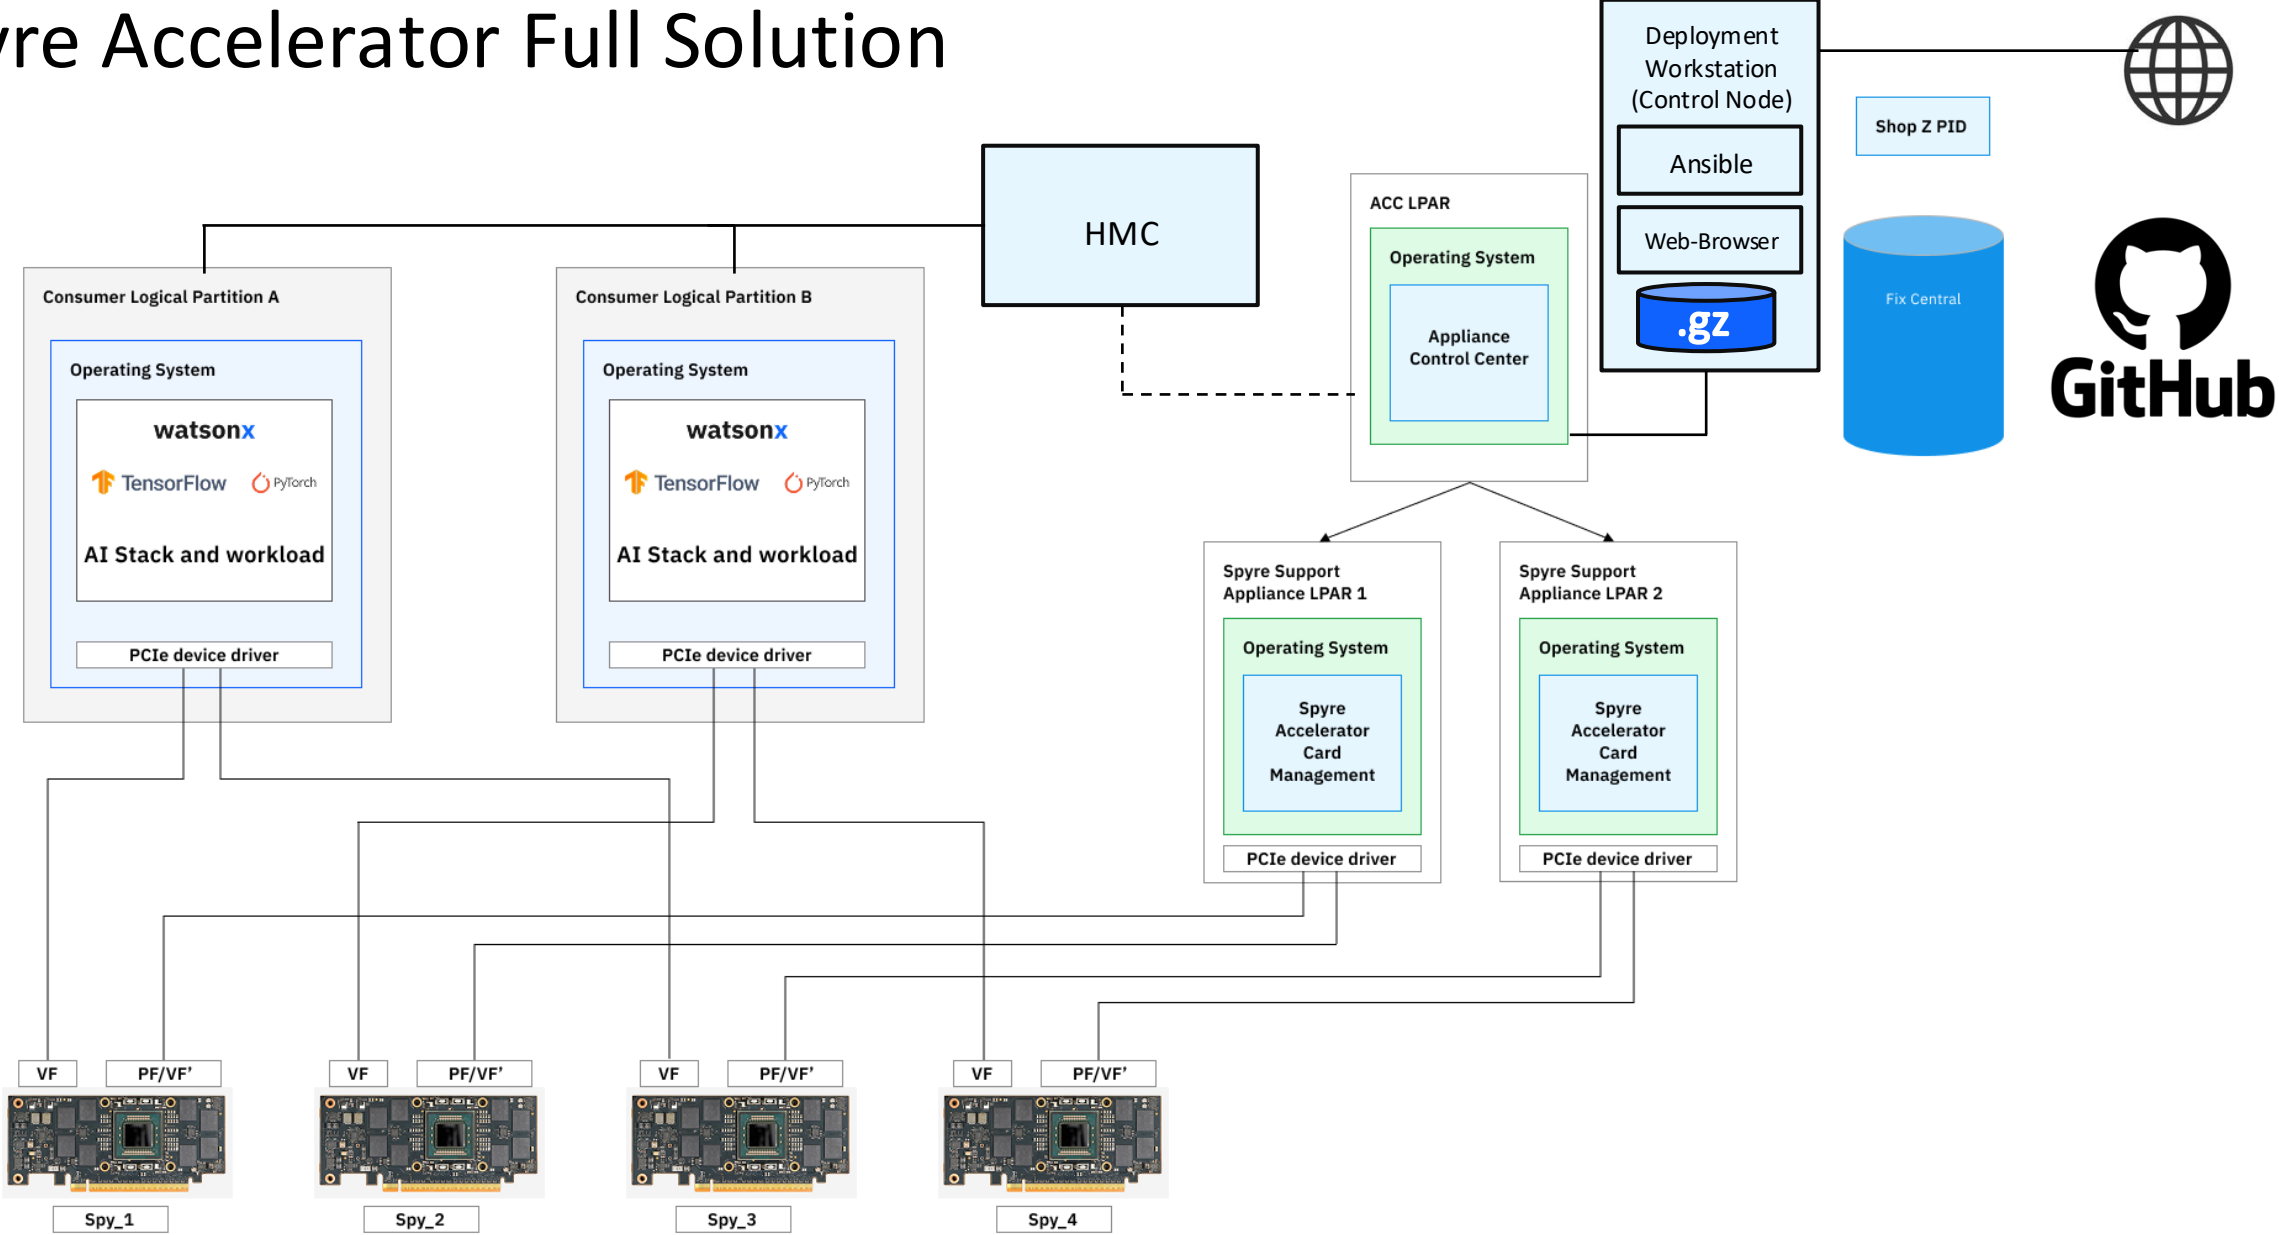


- Minimum 8 adapters
- Adapters are spread out across the fewest IO drawers due to power and cooling impacts
- Max of qty (8) AI accelerators in a single IO drawer
- Max of 48 total AI accelerators
- Qty (48) accelerators could be split in as little as 6 IO drawers

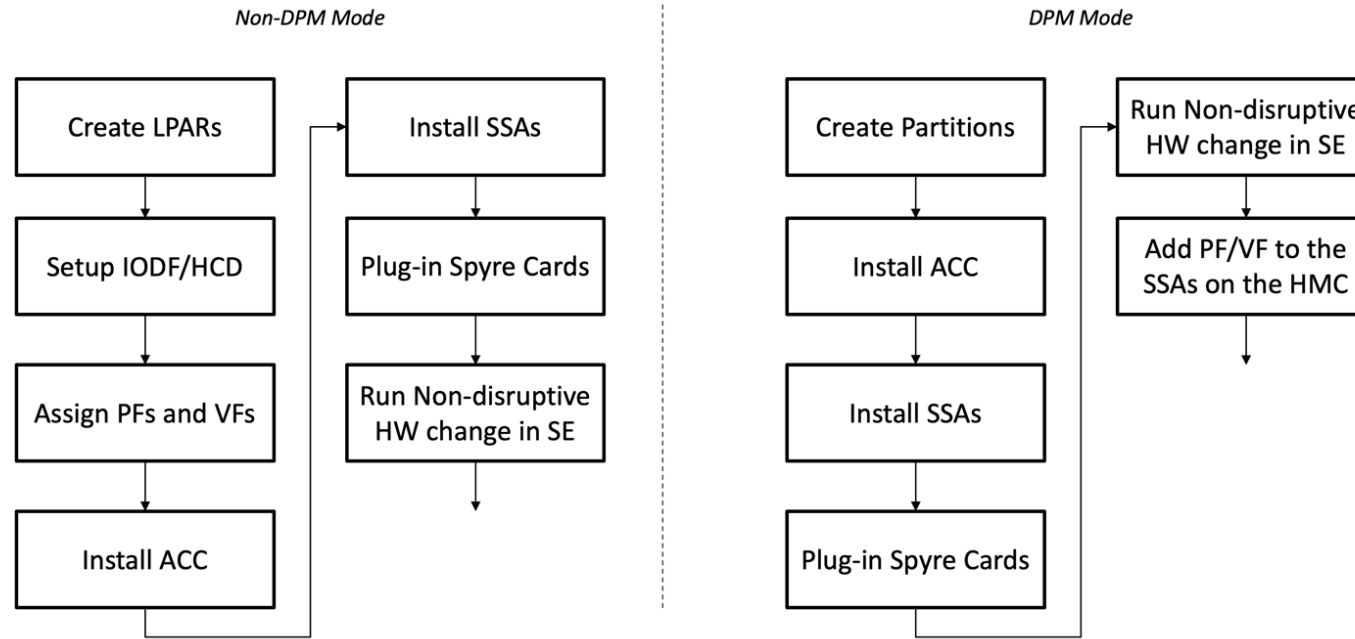
# Spyre System Overview



# Spyre Accelerator Full Solution

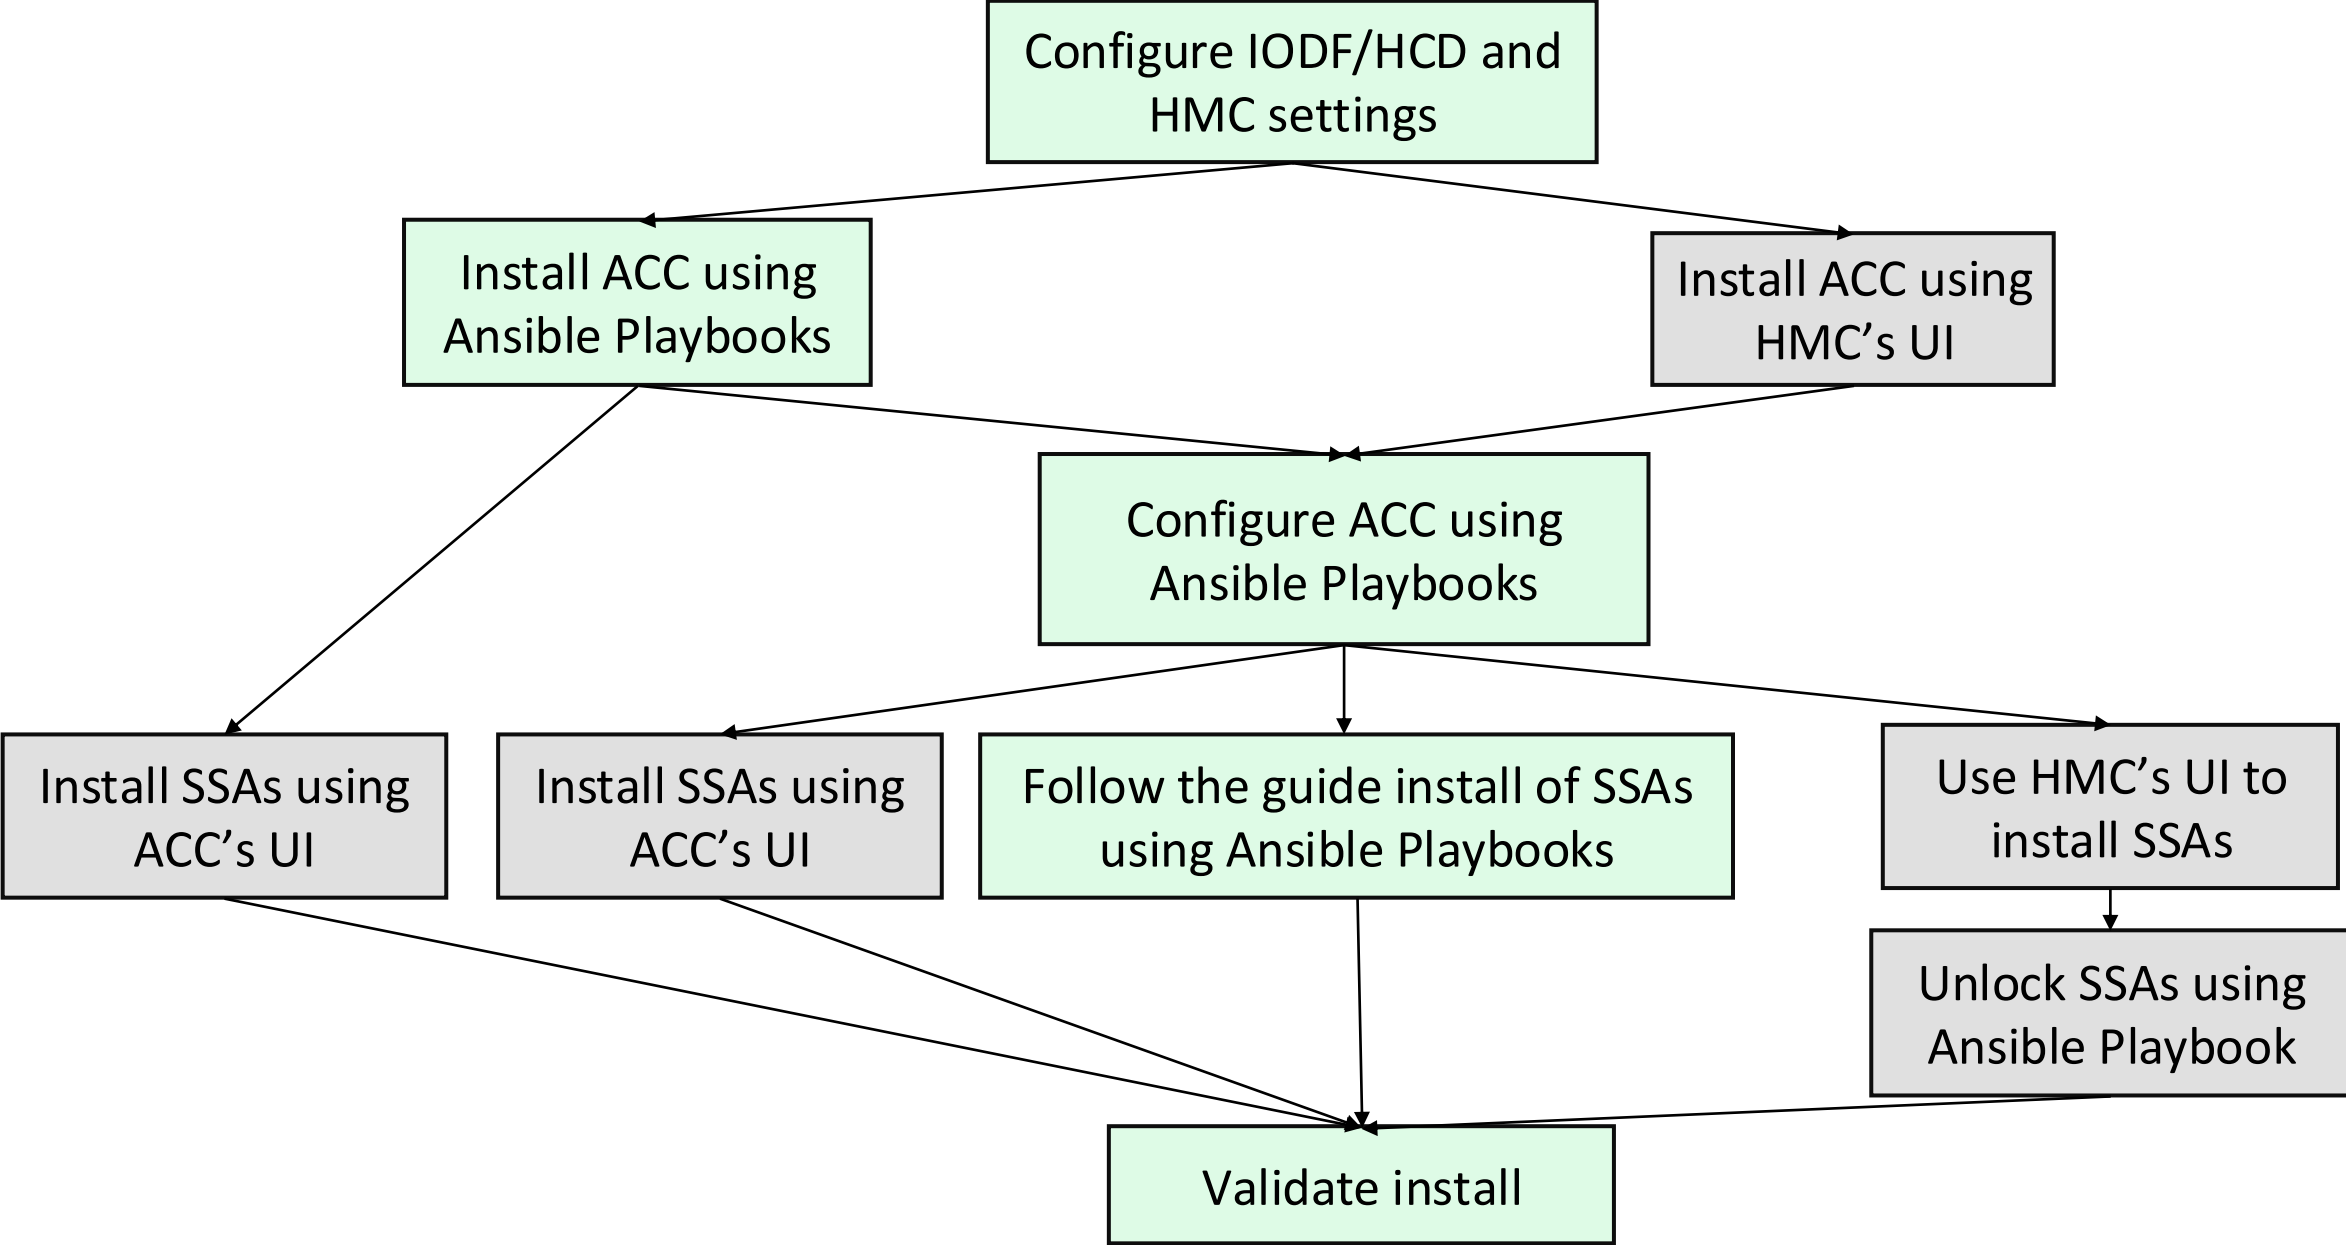


# Spyre Accelerator Workflow



1. Provision DASD and update HCD/IODF.
2. Provision three Secure Service Container (SSC) LPARs, two for Spyre Support Appliances (SSAs) and one for Appliance Control Center (ACC).
3. Define Physical Function (PF) and Virtual Functions (VF) as Function IDs (FIDs) per Spyre card, assign PF and first VF to SSA, and second VF to the consumer workload LPAR.
4. Plan and map Physical Channel ID (PCHIDs) for Spyre adapters using the Connectivity Mapping Tool.
5. Install ACC and SSA images using Hardware Management Console (HMC) and/or Support Element (SE) or ACC, then activate SSA LPARs.
6. Nondisruptive hot-plug and activation of Spyre cards – IBM System Service Representative (SSR) task

# Recommended Installation Path



# Requirements



# Hardware Requirements



## Hardware Requirements

- **3 SSC LPARs required**
  - **One Appliance Control Center**
    - 2 logical IFL/CP Shared
    - 16 GB Memory
    - 50 GB DASD (ECKD™/SCSI)
    - Networking connections
  - **Two Spyre Support Appliances**
    - 2 logical IFL/CP Shared each
    - 50 GB Memory each
    - 50 Gb DASD (ECKD/SCSI) each
    - Networking Connections for each

- Each system that has Spyre adapters will need **two Spyre Support Appliances** to manage the Spyre adapters
  - Recommended to split adapters evenly between the two SSCs
- Clients only need **one Appliance Control Center** in an environment as long as the SSC has network access to where they want to deploy Spyre Support Appliances
  - Acts as “App Store” to install the Spyre Support Appliance code
  - IBM requires at least 2 physical IFLs

# Software Requirements



## Software Requirements

- **Python 3.9+**
  - Used for ACC/SSA communication
- **Appliances**
  - ACC and SSA appliances download from IBM Fix Central
- **Ansible**
  - Used for ACC ansible playbooks

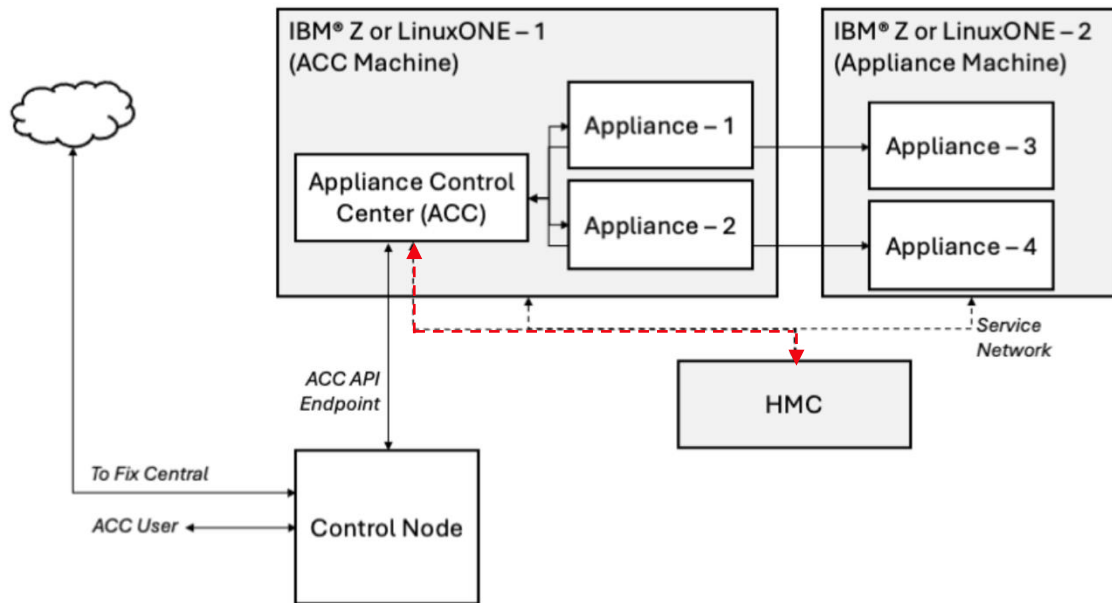
## Suggestions

- Review the Spyre Accelerator User guide for requirement details
- The control node must have Python & Ansible up and running
- HMC network connection to ACC is preferred and is the way to enable automation

# Choose installation mode

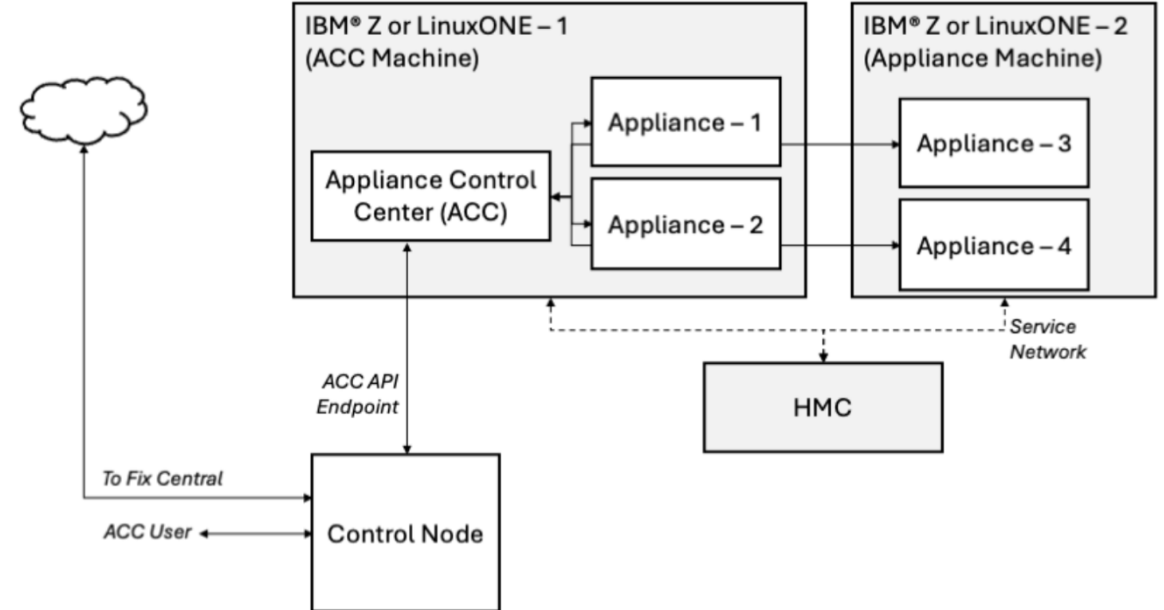
## Default Mode

- In the default mode, ACC has a network connectivity to HMC. This mode provides the maximum convenience and is suitable for automation.



## Standalone Mode

- In standalone mode, ACC does not have network connectivity to the HMC. As a result, certain actions on ACC must be preceded by manual steps that are performed on the HMC. This mode is suitable for environments where security, compliance, or organizational processes prevent ACC from communicating directly with the HMC.



# ACC Default mode

To run ACC in Default mode, you must ensure that the HMC connectivity is set up properly:

- ACC can communicate with HMC (both ways)
- Web services API must be enabled on the HMC
- Port 6794 (TCP) on the firewall between ACC and HMC must be opened
- ACC admin must have a user ID on the HMC and be authorized to edit activation profile, load, activate & deactivate LPARs
- No MFA on HMC login, in such cases, standalone mode ACC must be used.

If your network or your security rules don't allow you to meet all these needs, then choose Standalone mode,  
Else we recommend to use Default mode.

# Step by step walkthrough

- Hardware configuration
- ACC Setup
- SSA Setup
- Spyre cards installation



## Before you begin

- HCD User's Guide
- ACC and SSA must be on the same network
- ACC and SSA should be on CPCs that are managed by the same HMC

---

## About this task

For each OSA adapter used by the network adapter in SSC LPARs, you must define three device IDs in the I/O configuration:

- One device ID for network traffic
- Two device IDs for read/write control data

If the CHPID/FID is shared across multiple SSC LPARs, these three device addresses can also be shared by the SSC containers.

---

## Procedure

1. Define DASD
2. Create 3 LPARs in IODF
3. Create Physical Functions
4. Create Virtual Functions

# Define DASD

## Change Control Unit Definition

Specify or revise the following values.

Control unit number . . . . 0F50 +  
Control unit type . . . . . 2107 +  
Serial number . . . . . ABC01 +  
Description . . . . . DASD1 LCU 5D ACC SSA1 SSA2  
Connected to switches . . . 17 18 -- -- -- -- -- +  
Ports . . . . . DE A3 -- -- -- -- -- +  
Define more than eight ports . . 2 1. Yes  
2. No

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5= Reset  
F9=Swap      F12=Cancel



# Create LPARs

```
Processor ID . . . . . : B87          B87 z17
Configuration mode . . : LPAR
Channel Subsystem ID : 1

/ Partition Name  Number  Usage + UID  Description
- ACC            A       OS      N   ACC Appliance Control Center
- SSA1          C       OS      N   Spyre Support Appliance
- SSA2          D       OS      N   Spyre Support Appliance
```

```
View Partition / Channel Paths

Command ==> _____ Scroll ==> CSR

Processor ID . . . . . : B87          B87 z17
Configuration mode . . : LPAR
Partition name . . . . : ACC          ACC Appliance Control Center

Enter to continue.

CHPID   Type   Mode   Description
12      OSD   SPAN   Campus Network
AA      FC     SPAN   32G
B9      FC     SPAN   32G

***** Bottom of data *****

F1=Help   F2=Split   F3=Exit   F7=Backward   F8=Forward
F9=Swap   F12=Cancel F22=Command
```



# Create Physical Functions

```

Add PCIe Function

Specify or revise the following values.

Processor ID . . . . . : B87          B87 z17

Function ID . . . . . 4010
Type . . . . . PAIA          +

Channel ID . . . . . 18C +
Port . . . . . -          +
Virtual Function ID . . . . . +      Physical Function Y (Y/N)
Number of virtual functions 0__
UID . . . . . ---

Description . . . . . B87 SSA1 Physical Function
  
```

```

Define Access List

Command ==> _____ Scroll ==> CSR

Select one partition for the access list.

Function ID . . . . : 4010

/ CSS ID Partition Name Number Usage Description
/ 1 SSA1 C OS Spyre Support Appliance
- 1 SSA2 D OS Spyre Support Appliance
- 1 S5M 4 OS S5M - PLX5
- 1 S5V F OS S5V - PLX5
- 1 S51 1 OS S51 - PLX5
- 1 S53 3 OS S53 - PLX5
- 1 S55 5 OS S55 - PLX5
- 1 TA0 6 OS TA0 - PLXA
- 1 X5CFB87 7 CF X5CFB87 - PLX5
- 2 IESPLN1 5 OS zLinux - PyTorch Demo
- 2 R7A A OS R7A - PLX7
- 2 R72 2 OS R72 - PLX7

F1=Help F2=Split F3=Exit F5= Reset F6=Previous
F7=Backward F8= Forward F9=Swap F12=Cancel F22=Command
  
```



# Create Physical Functions

```
Define Candidate List

Command ==> _____ Scroll ==> CSR

Select one or more partitions for inclusion in the candidate list.

Function ID . . . . : 4010

/ CSS ID  Partition Name  Number  Usage  Description
- 1      R71                B       OS     R71 - PLX7
- 1      R75                2       OS     R75 - PLX7
- 1      R79                9       OS     R79 - PLX7
/ 1      SSA2              D       OS     Spyre Support Appliance
- 1      S5M                4       OS     S5M - PLX5
- 1      S5V                F       OS     S5V - PLX5
- 1      S51                1       OS     S51 - PLX5
- 1      S53                3       OS     S53 - PLX5
- 1      S55                5       OS     S55 - PLX5
- 1      TA0                6       OS     TA0 - PLXA
- 1      X5CFB87           7       CF     X5CFB87 - PLX5
- 2      IESPLN1           5       OS     zLinux - PyTorch

Demo

F1=Help      F2=Split      F3=Exit      F5= Reset    F6=Previous
F7=Backward  F8= Forward   F9=Swap      F12=Cancel   F22=Command
```



# Create Virtual Function 1

```
                Add PCIe Function
Specify or revise the following values.
Processor ID . . . . . : B87          B87 z17
Function ID . . . . . : 4011
Type . . . . . : PAIA          +
Channel ID . . . . . : 18C +
Port . . . . . : _ +
Virtual Function ID . . . . : 1__ +   Physical Function  N (Y/N)
Number of virtual functions 1
UID . . . . . : ---
Description . . . . . : B87 SSA1 1st Virtual Function
```

```
                Define Access List
Command ==> _____ Scroll ==> CSR
Select one partition for the access list.
Function ID . . . . : 4011
/ CSS ID Partition Name Number Usage Description
_ 1 LP1 E OS LP1 for COD
_ 1 R7H 8 OS R7H - PLX7
_ 1 R71 B OS R71 - PLX7
_ 1 R75 2 OS R75 - PLX7
_ 1 R79 9 OS R79 - PLX7
/ 1 SSA1 C OS Spyre Support Appliance
_ 1 SSA2 D OS Spyre Support Appliance
_ 1 S5M 4 OS S5M - PLX5
_ 1 S5V F OS S5V - PLX5
_ 1 S51 1 OS S51 - PLX5
_ 1 S53 3 OS S53 - PLX5
_ 1 S55 5 OS S55 - PLX5
F1=Help F2=Split F3=Exit F5= Reset F6=Previous
F7=Backward F8= Forward F9=Swap F12=Cancel F22=Command
```



# Create Virtual Function 1

```
Define Candidate List

Command ==> _____ Scroll ==> CSR

Select one or more partitions for inclusion in the candidate list.

Function ID . . . . : 4011

/ CSS ID Partition Name Number Usage Description
- 1 LP1 E OS LP1 for COD
- 1 R7H 8 OS R7H - PLX7
- 1 R71 B OS R71 - PLX7
- 1 R75 2 OS R75 - PLX7
- 1 R79 9 OS R79 - PLX7
/ 1 SSA2 D OS Spyre Support Appliance
- 1 S5M 4 OS S5M - PLX5
- 1 S5V F OS S5V - PLX5
- 1 S51 1 OS S51 - PLX5
- 1 S53 3 OS S53 - PLX5
- 1 S55 5 OS S55 - PLX5
- 1 TA0 6 OS TA0 - PLXA

F1=Help F2=Split F3=Exit F5= Reset F6=Previous
F7=Backward F8= Forward F9=Swap F12=Cancel F22=Command
```



# Create Virtual Function 2

```

Add PCIe Function

Specify or revise the following values.

Processor ID . . . . . : B87          B87 z17

Function ID . . . . . 4012
Type . . . . . PAIA          +

Channel ID . . . . . 18C +
Port . . . . .          +
Virtual Function ID . . . 2__ +   Physical Function  N (Y/N)
Number of virtual functions 1
UID . . . . .          ---

Description . . . . . B87 Linux lpar SPYDEMO
  
```

```

Define Access List

Command ==> _____ Scroll ==> CSR

Select one partition for the access list.

Function ID . . . . : 4012

/ CSS ID Partition Name Number Usage Description
- 4      R7D             4      OS      R7D - PLX5
- 4      R7J             E       OS      R7J - PLX7
/ 4      SPYDEMO         F       OS      zLinux - Spyre
- 4      S5A             A       OS      S5A - PLX5
- 4      S5D             D       OS      S5D - PLX5
- 4      S5G             7       OS      S5G - PLX5
- 4      S5J             2       OS      S5J - PLX5
- 4      S5N             8       OS      S5N - PLX5
- 4      S5S             B       OS      S5S - PLX5
- 4      TA3             3       OS      TA3 - PLXA
- 4      ZAIU            1       OS      ZAIU - LINUX
- 4      ZAIU1           6       OS      ZAIU1 - LINUX

F1=Help      F2=Split      F3=Exit      F5= Reset      F6=Previous
F7=Backward  F8= Forward   F9=Swap      F12=Cancel     F22=Command
  
```



# Summary

/	FID	CHID+	P+	VF+	PF	Type+	UID	Description
-	4010	18C	-	--	Y	PAIA	----	B87 SSA1 Physical Function
-	4011	18C	-	1	N	PAIA	----	B87 SSA1 1st Virtual Function
-	4012	18C	-	2	N	PAIA	----	B87 Linux lpar SPYDEMO

```
FUNCTION FID=4010,PF,PCHID=18C,PART=((SSA1),(SSA2)),TYPE=PAIA
FUNCTION FID=4011,VF=1,PCHID=18C,PART=((SSA1),
(SSA2)),TYPE=PAIA
FUNCTION FID=4012,VF=2,PCHID=18C,PART=((SPYDEMO),(SSA2,S5J,ZAIU,ZAIU1)),TYPE=PAIA
```

# Appliance Control Center Setup



## Before you begin

- Detailed [ACC User's Guide](#)
- Know your IBMid
- The target FCP disks and FICON DASD must be large enough to fit the uncompressed appliance.

---

## About this task

- You must configure and start a Secure Service Container partition with the boot option SSC Installer selected
- Identify the IP address for the Secure Service Container partition
- Order and download the ACC and SSA to your local disk.
- To avoid network latency, upload the appliance image from a browser on the same network as the Installer
- If other workloads are consuming bandwidth on the storage server or channel subsystem, consider performing the installation when the system is less busy

---

## Procedure

1. Fix Central login with IBMid
2. Find "IBM Z Appliance Control Center"
3. Browse for fixes and download the tar file
4. Untar the file using "tar -xvf fix\_name.tar"
5. Update Activation Profile

# Downloading Fix

## Fix Central

[Help](#)

Search all IBM for Fix metadata

Filter search by - IBM Z Appliance Control Center

### Select fixes

Other Software, IBM Z Appliance Control Center (1.2.11, Linux zSeries)

**Need to download your product?**

[Find full product install images on Passport Advantage](#)

[Show fix details](#) | [Hide fix details](#)

<input type="checkbox"/>	Description	Release date
<input checked="" type="checkbox"/>	<p><b>1</b> fix pack: → <a href="#">1.2.11-SystemZFamily-ACC-GA-release</a> Appliance Control Center for IBM Z &amp; IBM LinuxONE 1.2.11</p> <p><b>Platforms:</b> All</p> <p><b>Applies to versions:</b> 1.2.11</p> <p><b>Upgrades to:</b> 1.2.11</p> <p><b>Severity:</b></p> <p><b>Categories:</b></p> <p><b>Abstract:</b> Appliance Control Center for IBM Z &amp; IBM LinuxONE 1.2.11</p> <p><b>Restrictions:</b> license</p> <p><a href="#">More Information</a></p>	2025/12/16

1-1 of 1 results

[Show fix details](#) | [Hide fix details](#)

#### Download options

[Share this download list](#)

#### Change your selection

Product selector  
IBM Z Appliance Control Cent

Installed Version  
1.2.11

Platform  
Linux zSeries

#### Filter your content

**Fix type**  
 fix pack (1)

**And Applies to**  
 1.2.11 (1)

**And Platform**



# Update Activation Profile

The screenshot shows the HMC Management console with a table of SSCs. A context menu is open for SSC12, and the 'Operational Customization' option is selected, which has opened a sub-menu with 'Customize/Delete Activation Profiles' highlighted.

SSC ID	Status	SSC Name	SSC Name
SSC11	Not activated	SSC11	SSC11
SSC12	Not activated	SSC12	SSC12
SSC13	Not activated	SSC13	SSC13
SSC14	Not activated	SSC14	SSC14
SSC15	Not activated	SSC15	SSC15
SSC16	Not activated	SSC16	SSC16

The dialog box shows a table of activation profiles. The 'SSC12' profile is selected. The table has columns for 'Select', 'Name', 'Type', and 'Profile Description'. Below the table, it shows 'Total: 2 Filtered: 2 Selected: 1' and buttons for 'Customize profile', 'Delete', 'Close', and 'Help'.

Select	Name	Type	Profile Description
<input checked="" type="radio"/>	SSC12	Image	Test using python-zhmcclient
<input type="radio"/>	DEFAULTLOAD	Load	This is the default Load profile.

Total: 2 Filtered: 2 Selected: 1

Buttons: Customize profile, Delete, Close, Help



# Update Activation Profile

This screenshot shows the 'Identify Profile' step of the 'New Image Profile' configuration. The left sidebar lists the steps: Welcome, Identify Profile (current), Select Operating System, Assign Processors, Enable Performance Controls, Specify Processing Weights, Allocate Storage, Configure Load, Select Crypto Domains, Select Crypto Candidate, Activate Crypto Online, Assign SSC Adapters, and Summary. The main area contains instructions: 'Specify a hexadecimal value to use as the partition identifiers used by the logical partition. The partition identifier must be unique among the identifiers of other active logical partitions.' and 'Also specify a brief note (max 50 characters), that describes the contents or purpose of the profile.' The 'Partition Identifier' field contains '1C' and the 'Profile Description' field is empty. Under 'Initialized to:', the 'Use factory defaults' radio button is selected. A 'Profile' dropdown menu is set to 'DEFAULT'. Navigation buttons at the bottom include '< Back', 'Next >', 'Finish', and 'Cancel'.

This screenshot shows the 'Allocate Storage' step of the 'New Image Profile' configuration. The left sidebar shows the progress: Welcome, Identify Profile, Select Operating System, Assign Processors, Enable Performance Controls, Specify Processing Weights, and Allocate Storage (current). The main area contains instructions: 'Central storage is allocated when this partition is initialized. This storage not shared with other active logical partitions.' and 'Enter the amount of central storage to allocate to the logical partition upon activation.' It also states: 'Central storage range is from 1024MB to 8224GB. The central storage amount must be a multiple of 1024MB.' The 'Storage' field is set to '50' and the 'Units' dropdown is set to 'Gigabytes (GB)'. Below this, the 'Virtual Flash Memory (GB)' section is visible, with instructions: 'Choose memory amounts in 16GB increments up to a maximum of 1536GB.' The 'Initial' and 'Maximum' fields are both set to '0' with up/down arrow controls. Navigation buttons at the bottom include '< Back', 'Next >', 'Finish', and 'Cancel'.



# Update Activation Profile

The screenshot displays the IBM Z activation profile configuration interface. The main window is titled "New Image Profile - Image profile: SSA1 - B87" and shows a progress list on the left with "Assign SSC Adapters" selected. The "Assign SSC Adapters" section includes fields for Administrator user ID (admin), Administrator password, Confirm password, and Host name (b87-ssa1). Below this is a "Network Adapters" table with columns for Select, CHPID, FID, Port, VLAN, IP address, and Mask/Prefix. A table with one row is shown: Select (radio), CHPID (12), FID, Port (1), VLAN (2231), IP address (9. .151), Mask/Prefix (20). Below the table are fields for IPv4 gateway (9. .254) and IPv6 gateway. A "DNS Servers" section is also present with a table for Select, IP address, and a "Select Action" dropdown.

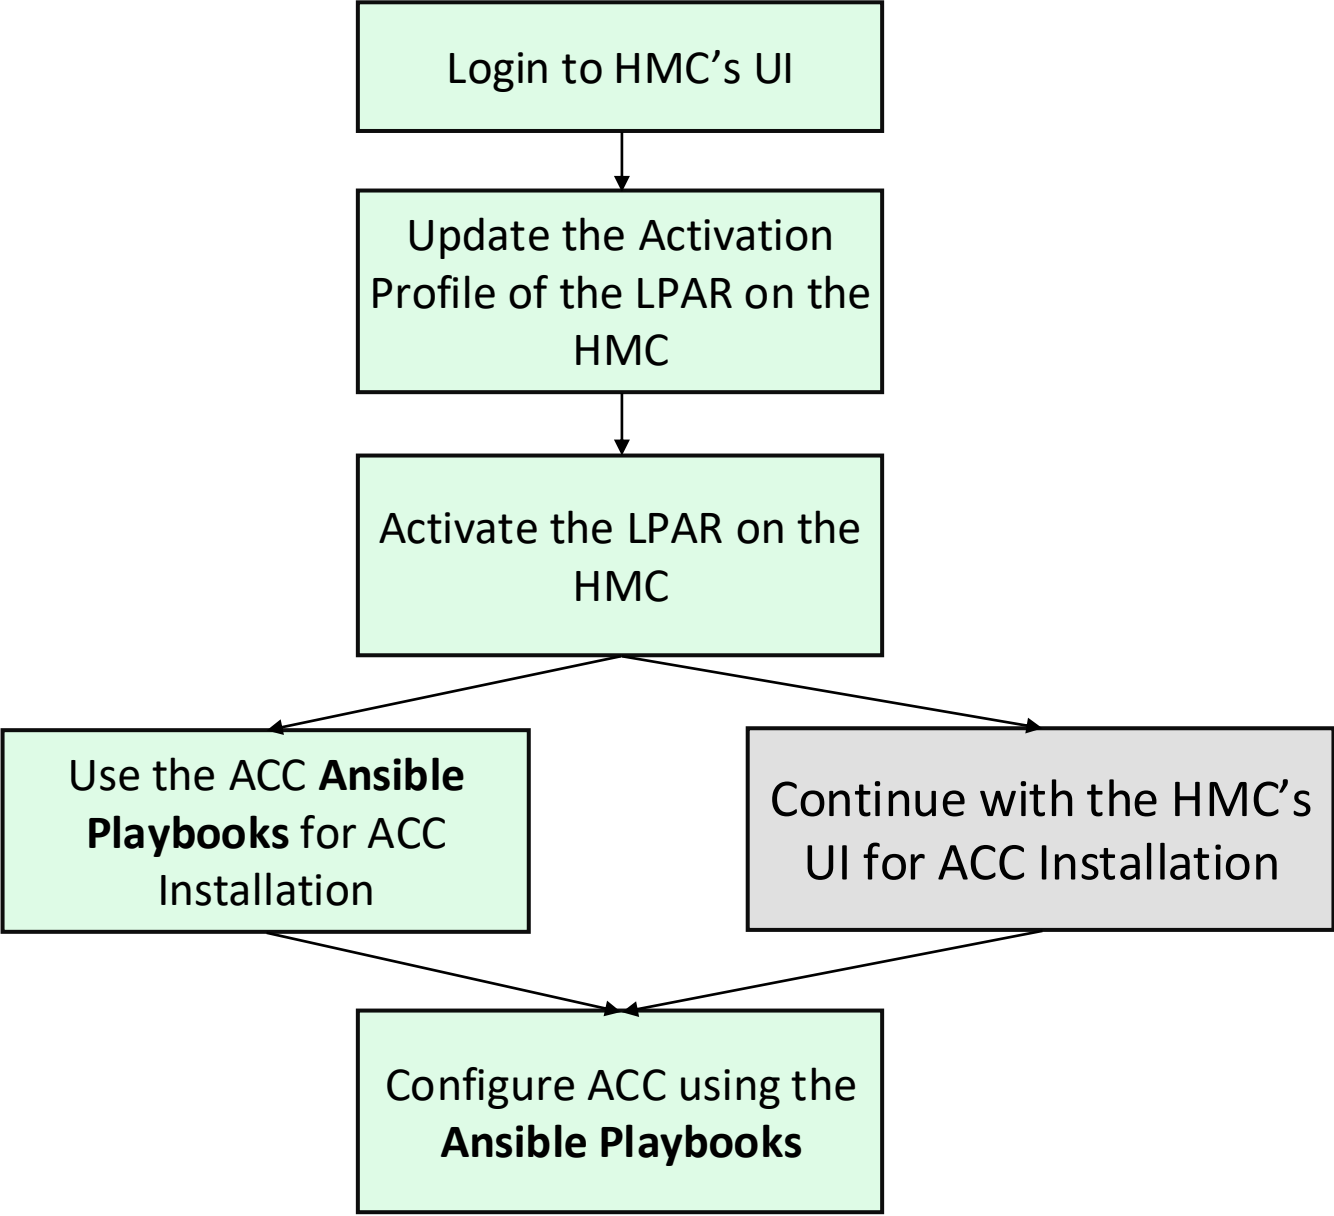
An "Add Network Adapter - B87" dialog box is open in the foreground, titled "Add Network Adapter - B87". It contains the following fields and options:

- Adapter Type:  CHPID  FID
- CHPID: 12
- Port:  0  1
- VLAN ID (1 - 4094): 2231
- IP address type:  DHCP  DHCP IPv6  Static IPv4  Static IPv6
- IP address: \* 9. .151
- Mask: \* 20

Buttons for "OK", "Cancel", and "Help" are at the bottom of the dialog.



# Install ACC



# Install ACC

## Before you begin

- Download the Appliance Control Center for IBM Z & IBM LinuxONE installation image from the Fix Central to the control node
- Download and save the scripts that are available on [GitHub](#)
- Ensure that the LPARs are created and available on the HMC.
- Ensure that you have the SSC network information
- Ensure that Ansible is already installed on the control node, for example, a laptop

---

## About this task

- We recommend to use Ansible Playbooks for ACC Installation to save time and less user error than HMC's UI

---

## Procedure

1. Clone Github
2. Setup Python environment
3. Export HMC credentials
4. Modify the variables in `appliance_deploy_default_ansible/acc_env_vars.yaml` to match your environment
5. Install ACC using playbook `ansible-playbook ./appliance_deploy_default_ansible/00_acc_install.yaml`

# ACC Default mode

## 1. Clone the git repository

```
Cloning into 'z_ansible_collections_samples'...
remote: Enumerating objects: 18160, done.
remote: Counting objects: 100% (2625/2625), done.
remote: Compressing objects: 100% (551/551), done.
remote: Total 18160 (delta 2258), reused 2080 (delta 2073), pack-reused 15535 (from 3)
Receiving objects: 100% (18160/18160), 109.39 MiB | 3.70 MiB/s, done.
Resolving deltas: 100% (9505/9505), done.
Updating files: 100% (1995/1995), done.
```

## 2. Setup Python environment

1. Create a virtual environment in acc\_install\_ansible folder

2. Install packages

```
pip install click click_shell zhmcclient urllib3
```

3. Exit the virtual environment

## 3. Export your HMC credentials by running the command:

```
export HMC_USER=<your_HMC_username>
```

```
export HMC_PASSWORD=<your_HMC_password>
```

# ACC Default mode

4. Modify the variables in `appliance_deploy_default_ansible/acc_env_vars.yaml` to install ACC
5. To install ACC, run the following playbook:  
`ansible-playbook ./appliance_deploy_default_ansible/00_acc_install.yaml`
6. Configure AAC to manage SSA LPARs (Ansible method)

File	Purpose
<code>admin_vars.yaml</code>	This file contains variables used by the ACC administrator to initialize ACC and assign resources to the appliance owners
<code>owner_vars.yaml</code>	This file contains variables used by the appliance owners to install their appliances
<code>01_admin_actions.yaml</code>	ACC administrator can use this playbook to initialize ACC
<code>02a_assign_1_lpar.yaml</code>	ACC administrator can use this playbook to assign a single LPAR to an appliance owner
<code>02b_assign_2_lpar.yaml</code>	ACC administrator can use this playbook to assign two LPARs to an appliance owner
<code>03_owner_actions.yaml</code>	Appliance owner can use this playbook to initialize the user on ACC
<code>04_install_flow.yaml</code>	Appliance owner can use this playbook to install and activate an appliance

# ACC Default mode

File	Purpose
<code>admin_vars.yaml</code>	This file contains variables used by the ACC administrator to initialize ACC and assign resources to the appliance owners
<code>01_admin_actions.yaml</code>	ACC administrator can use this playbook to initialize ACC
<code>02b_assign_2_lpar.yaml</code>	ACC administrator can use this playbook to assign two LPARs to an appliance owner

## 6. Configure AAC to manage SSA LPARs (Ansible method)

Navigate to `appliance_deploy_ansible` directory and update the variables present in `admin_vars.yaml` file.

`cc_url` - ACC IP address.

`z_machine_lpar` - is the LPAR on which the appliance is installed.

`app_ip` - LPAR IP address.

Any additional variables related to SSA LPARs.

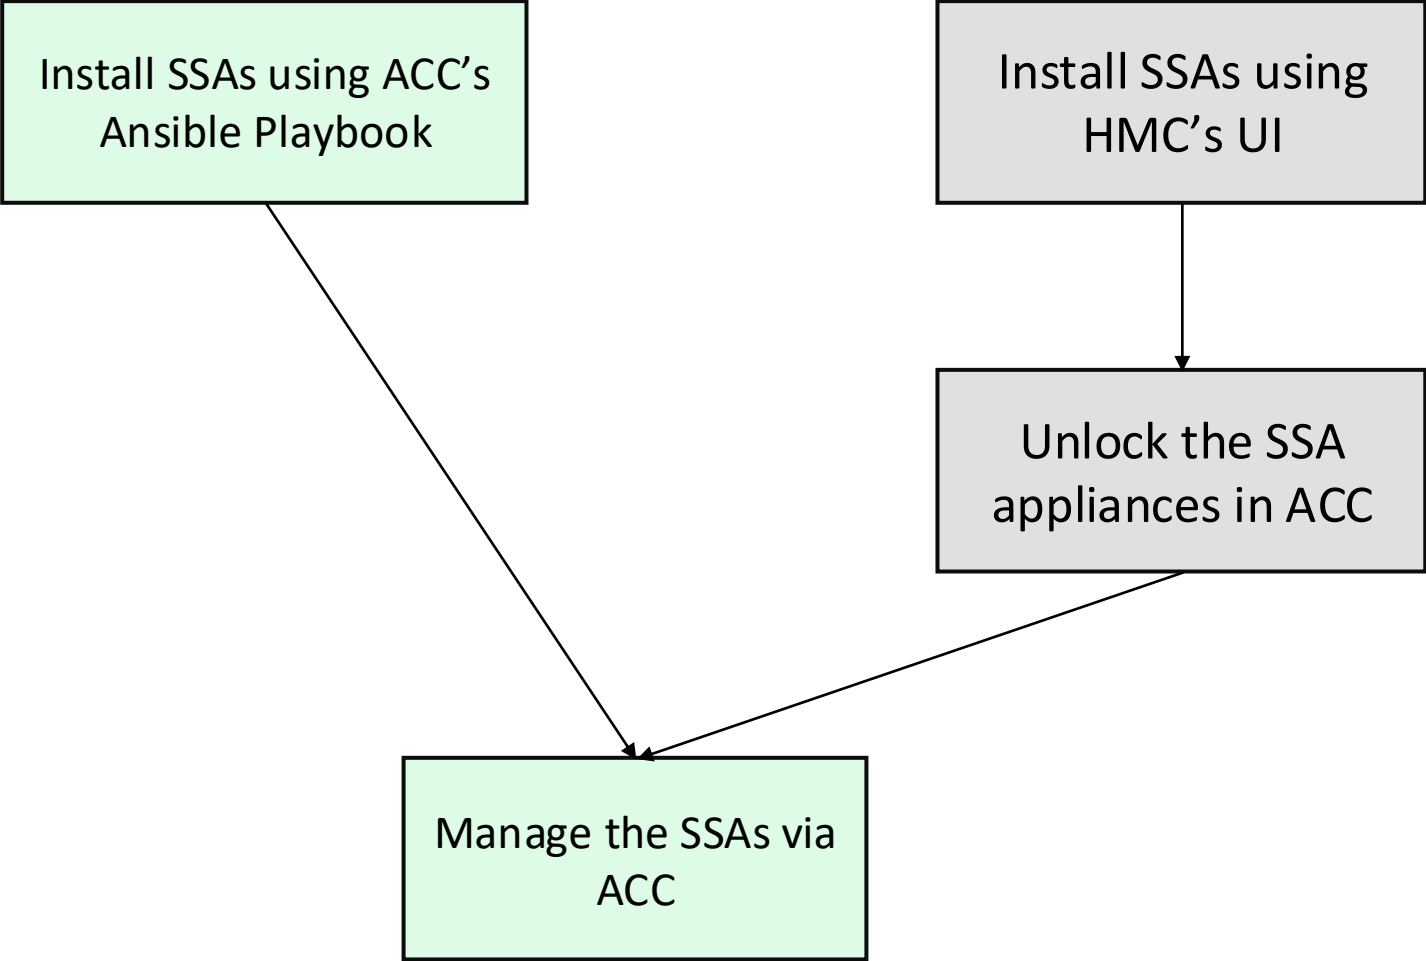
Initialize and configure ACC by running:

```
ansible-playbook 01_admin_actions.yaml
```

Configure the SSA LPARs to be a part of a resource group by running:

```
ansible-playbook 02b_assign_2_lpar.yaml
```

# Installing SSA



## Before you begin

---

- Detailed [SSA User's Guide](#)
- Know your IBMid
- The target FCP disks and FICON DASD must be large enough to fit the uncompressed appliance.

## About this task

- You must configure and start a Secure Service Container partition with the boot option SSC Installer selected
- Identify the IP address for the Secure Service Container partition
- Order and download the SSA to your local disk.
- To avoid network latency, upload the appliance image from a browser on the same network as the Installer
- If other workloads are consuming bandwidth on the storage server or channel subsystem, consider performing the installation when the system is less busy

## Procedure

---

1. Fix Central
2. Install using Ansible or ACC UI

# Downloading Fix

**Fix Central**  
Help

Search all IBM for Fix metadata

Filter search by - IBM Z Spyre Support Appliance

## Select fixes

Other Software, IBM Z Spyre Support Appliance (1.1.65, Linux zSeries)

### Need to download your product?

[Find full product install images on Passport Advantage](#)

**Continue**  **Clear selections** [Show fix details](#) | [Hide fix details](#)

<input type="checkbox"/>	<input type="checkbox"/> <b>Description</b>	<input type="checkbox"/> <b>Release date</b>
<input type="checkbox"/>	<p><b>1</b> fix pack: → <a href="#">1.1.65-SystemZFamily-SSA-GA-release</a> Spyre Support Appliance for IBM Z &amp; IBM LinuxONE 1.1.65</p> <p><b>Platforms:</b> All</p> <p><b>Applies to versions:</b> 1.1.65</p> <p><b>Upgrades to:</b> 1.1.65</p> <p><b>Severity:</b></p> <p><b>Categories:</b></p> <p><b>Abstract:</b> Spyre Support Appliance for IBM Z &amp; IBM LinuxONE 1.1.65</p> <p><b>Restrictions:</b> license</p> <p><a href="#">More Information</a></p>	2025/12/11

1-1 of 1 results

**Continue**  **Clear selections**  **Back** [Show fix details](#) | [Hide fix details](#)

**Download options**

[Share this download list](#)

**Change your selection**

Product selector

Installed Version

Platform

**Filter your content**

**Fix type**  
 fix pack (1)

**And Applies to**  
 1.1.65 (1)

**And Platform**



# ACC Default mode

File	Purpose
<code>owner_vars.yaml</code>	This file contains variables used by the appliance owners to install their appliances
<code>03_owner_actions.yaml</code>	Appliance owner can use this playbook to initialize the user on ACC
<code>04_install_flow.yaml</code>	Appliance owner can use this playbook to install and activate an appliance

## Update configuration files

- Modify the variables in the `owner_vars.yaml` file.
- Update the `image_path` variable in the `owner_vars.yaml` file to specify the location of the downloaded image.

Run the following playbook to update owner credentials and upload the appliance image:

```
ansible-playbook 03_owner_actions.yaml
```

Configure the SSA LPARs to be a part of a resource group by running:

```
ansible-playbook 04_install_flow.yaml
```

# Documentation



Official documentation published by IBM include:

- [IBM Spyre Accelerator User's Guide \(GC28-7071-00\)](#)
- [Appliance Control Center for IBM Z & IBM LinuxONE User's Guide \(GC28-7073-00\)](#)
- [Spyre Support Appliance for IBM Z & IBM LinuxONE User's Guide \(GC28-7072-00\)](#)

Official resource:

- [IBM Spyre support page](#)
- [Fix Central](#)
- [Ansible playbooks for ACC](#)

# Statements of Direction

**Red Hat OpenShift AI:** IBM and Red Hat intend to support deployment of Red Hat OpenShift AI across IBM® LinuxONE and IBM Z (via Linux on IBM Z, and z/OS® Container Extensions (zCX)) for a broad spectrum of options to deploy and manage the full lifecycle of predictive and generative AI models at scale.

**Red Hat AI Inference Server:** IBM and Red Hat intend to support Red Hat AI Inference Server on IBM® LinuxONE and IBM Z (via Linux on IBM Z and z/OS Container Extensions (zCX)) to deliver vLLM benefits for optimized generative AI. As part of z17 and IBM® LinuxONE 5, the Spyre Accelerator is designed for vLLM and LLM usage, enabling on-premises integration of AI solutions on a platform designed for mission-critical performance.

**zCX Support:** IBM intends to support the deployment of watsonx Assistant for IBM Z on zCX. These applications will leverage the Spyre Accelerator to deliver enhanced performance, scalability, and efficiency for AI-powered development and operations on IBM Z.

**AI Toolkit for IBM Z & IBM® LinuxONE with Spyre:** IBM intends to deliver capabilities to exploit the Spyre Accelerator with components of the AI Toolkit for IBM Z & IBM® LinuxONE.

**Machine Learning for IBM z/OS with Spyre:** IBM intends to deliver capabilities to exploit the Spyre Accelerator with Machine Learning for IBM z/OS.



# Your feedback is important!

## Submit a session evaluation for each session you attend:

[www.share.org/evaluation](http://www.share.org/evaluation)

