

Technical Advisory Council Meeting

July 16, 2020

 THE **LINUX** FOUNDATION

 LF AI

Antitrust Policy Notice

- › Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws.
- › Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at <http://www.linuxfoundation.org/antitrust-policy>. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrave of the firm of Gesmer Undergone LLP, which provides legal counsel to the Linux Foundation.

Recording of Calls

Reminder:

TAC calls are recorded and available for viewing on the [TAC Wiki](#)

Reminder: LF AI Useful Links

Web site: lfai.foundation
Wiki: wiki.lfai.foundation
GitHub: github.com/lfai
Landscape: landscape.lfai.foundation or l.lfai.foundation
Mail Lists: <https://lists.lfai.foundation>

LF AI Logos: <https://github.com/lfai/artwork/tree/master/lfai>

LF AI Presentation Template:

https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk_-czASlz2GTBRZk2/view?usp=sharing

Events Page on LF AI Website: <https://lfai.foundation/events/>

Events Calendar on LF AI Wiki (subscribe available):

<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091544>

Event Wiki Pages: <https://wiki.lfai.foundation/display/DL/LF+AI+Foundation+Events>

Agenda

- › Roll Call
- › Approval of Minutes
- › Guest Presentation: Mindspore
- › LF AI General Updates
- › Upcoming TAC Meetings
- › Open Discussion

TAC Voting Members

Member	Contact	Email
AT&T	Reuben Klein	rk1518@att.com
Baidu	Daxiang Dong	dongdaxiang@baidu.com
Ericsson	Rani Yadav-Ranjan	rani.yadav-ranjan@ericsson.com
Huawei	Huang Zhipeng	huangzhipeng@huawei.com
Nokia	Pantelis Monogioudis	pantelis.monogioudis@nokia.com
Tech Mahindra	Nikunj Nirmal	nn006444@techmahindra.com
Tencent	Bruce Tao	brucetao@tencent.com
Zilliz	Jun Gu	jun.gu@zilliz.com
ZTE	Wei Meng	meng.wei2@zte.com.cn
Acumos AI Project	Nat Subramanian	natarajan.subramanian@techmahindra.com
Angel Project	Bruce Tao	brucetao@tencent.com
ONNX Project	Jim Spohrer*	spohrer@us.ibm.com

* TAC Chairperson

Approval of Minutes

Draft minutes from the June 18th meeting of the TAC were previously distributed to the TAC members

Proposed Resolution:

- › That the minutes of the June 18th meeting of the Technical Advisory Council of the LF AI Foundation are hereby approved

Guest Presentation: Mindspore

Zhipeng (Howard) Huang

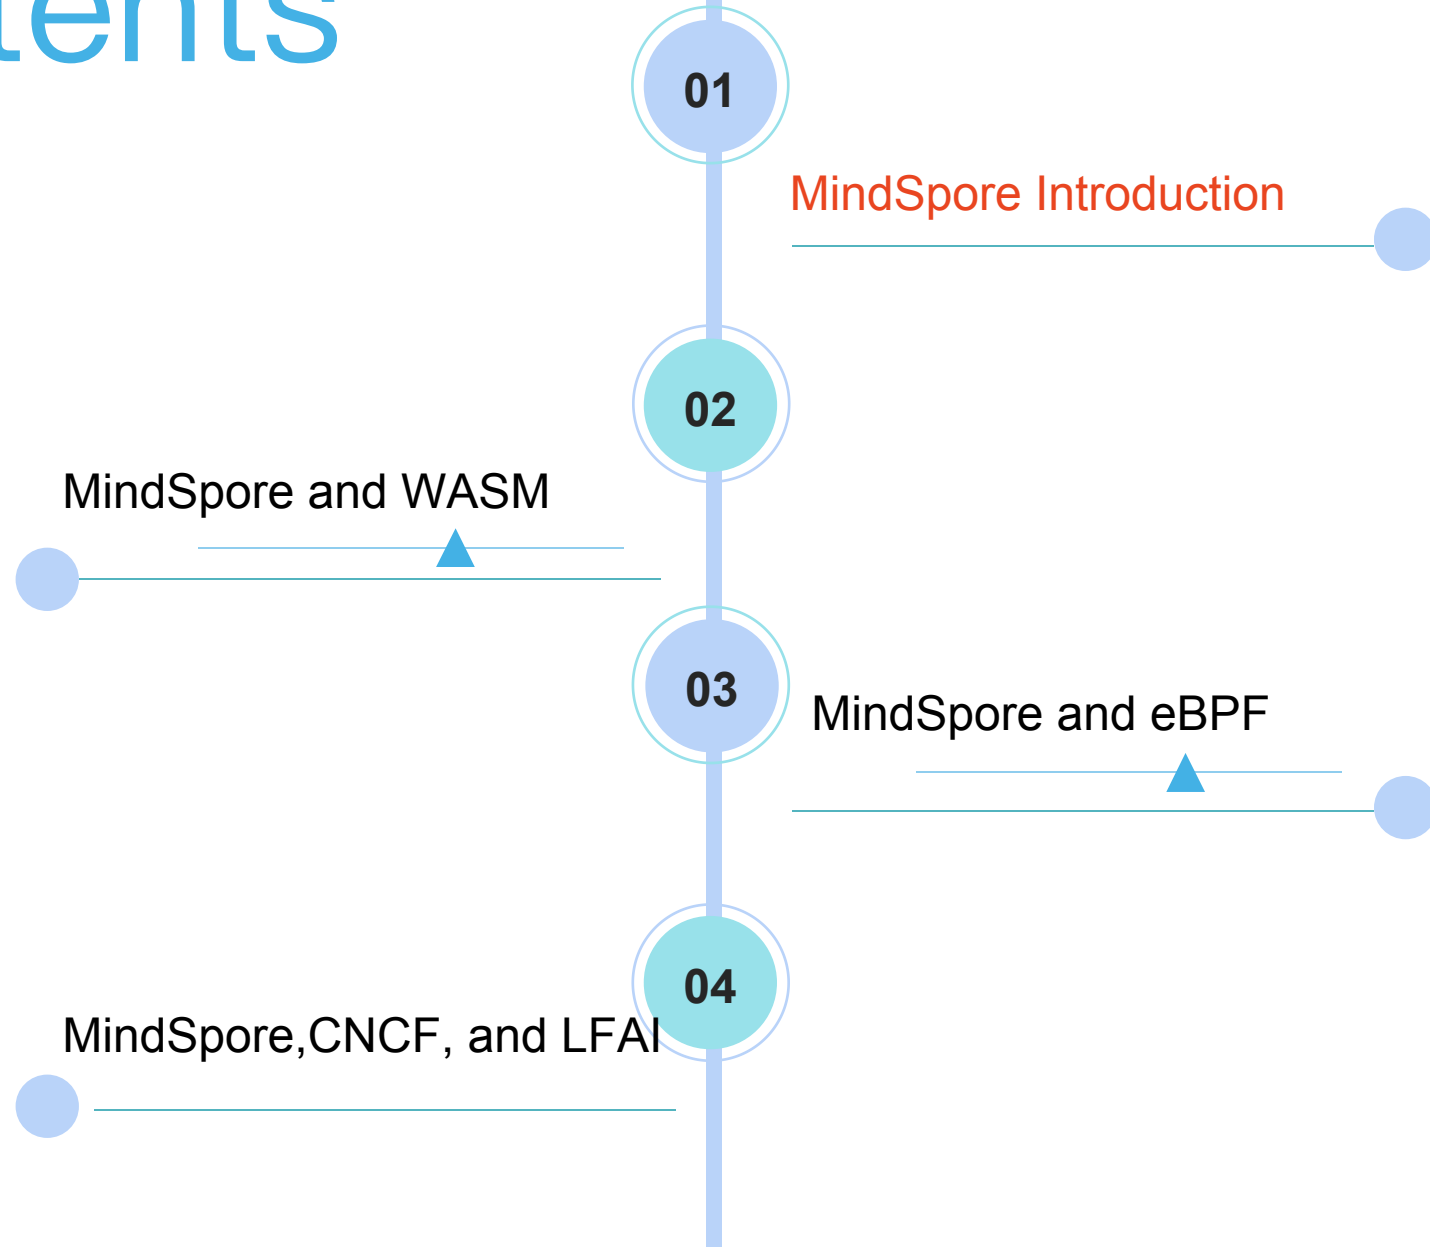


MindSpore

MindSpore Introduction

Zhipeng Huang

Contents



01

MindSpore Introduction

02

MindSpore and WASM

03

MindSpore and eBPF

04

MindSpore, CNCF, and LFAI

MindSpore Introduction



MindSpore is a new open source deep learning training/inference framework that could be used for mobile, edge and cloud scenarios.

MindSpore is designed to provide development experience with friendly design and efficient execution for the data scientists and algorithmic engineers, native support for Ascend AI processor, and software hardware co-optimization.

At the meantime MindSpore as a global AI open source community, aims to further advance the development and enrichment of the AI software/hardware application ecosystem.



<https://www.mindspore.cn>

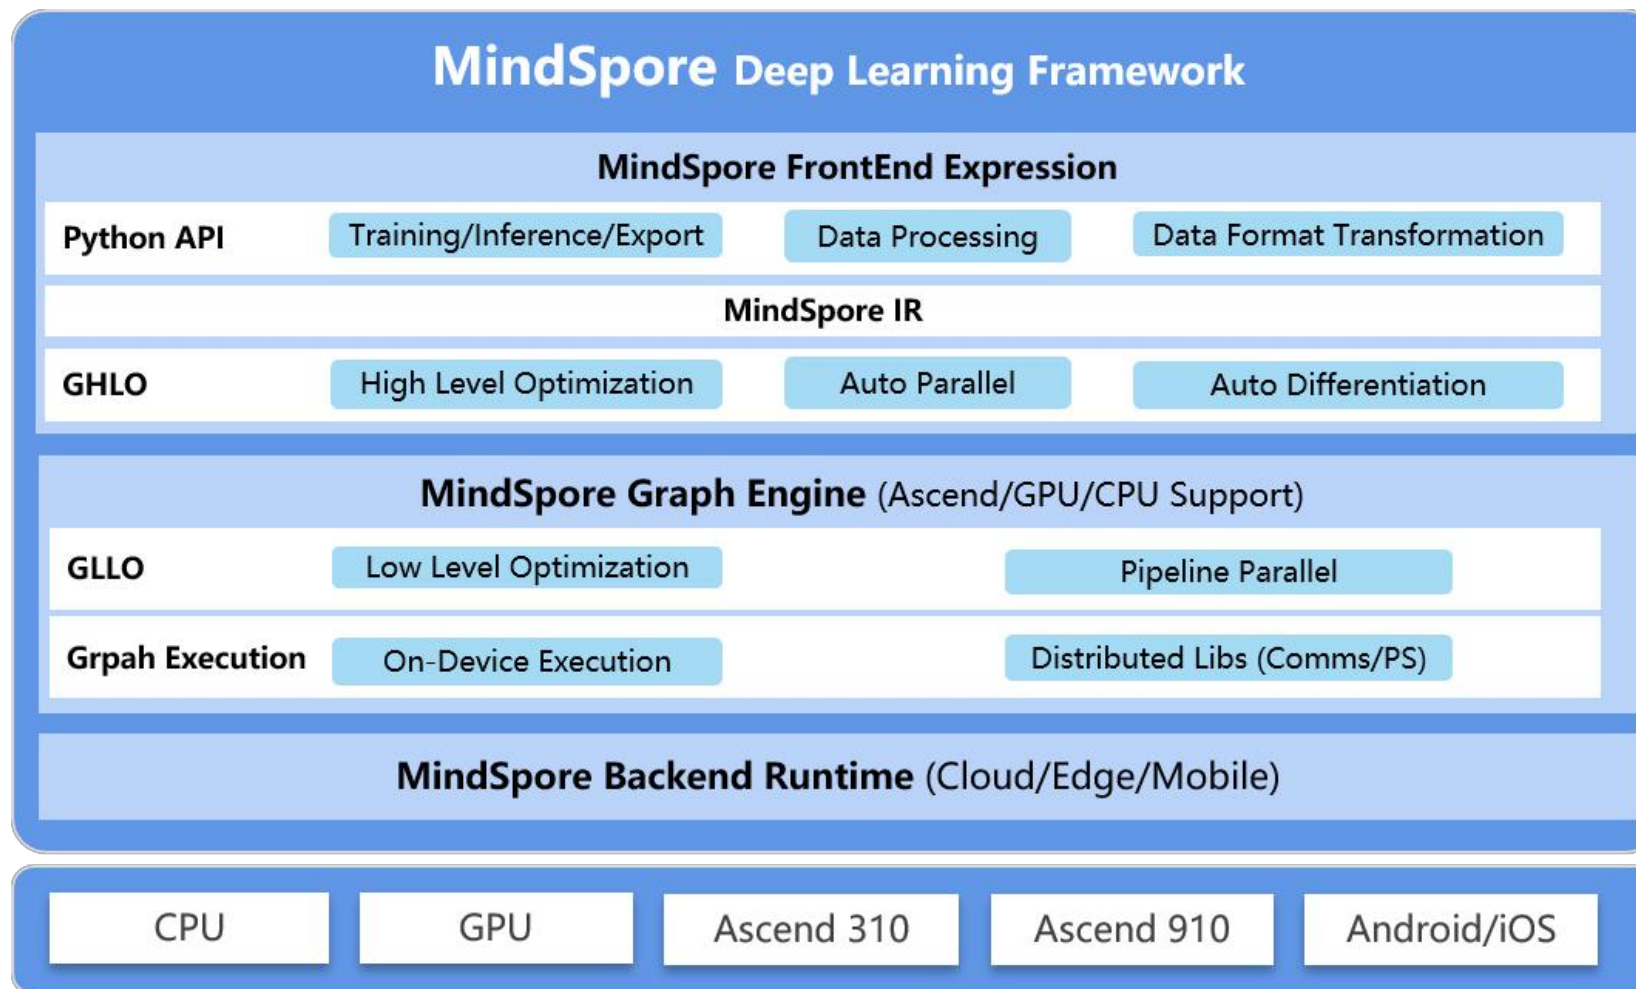


<https://www.gitee.com/mindspore>





<https://www.github.com/mindspore-ai>


Overview



Technical Steering Committee



14 members from various universities/institutions/companies that forms an open and global technical governing body



SIG/WG



MindSpore

SIGs

WGs

FrontEnd,
Compiler
Executor
ModelZoo
Data
GraphEngine
Visualization
Security
...

Documentation
Infrastructure
...

mindspore-ai / community


Unwatch 5 Star 4 Fork 0

Code Issues 0 Pull requests 0 Actions Projects 0 Wiki Security Insights Settings



Branch: master community / sigs / Create new file Upload files Find file History

Commit	Message	Time
dengyiping2014	add frontend sig	Latest commit becc448 5 days ago
..	..	
data	update notes for data sig 001-20200402.md	6 days ago
executor	add executor sig	6 days ago
frontend	add frontend sig	5 days ago
visualization	add visualization repo for sig	7 days ago

Open Development



Community partners



Open Collaboration

Auto Differentiation




MindSpore

Interface layer




$$y(n) = \sum_{i=-\infty}^{\infty} x(i)h(n-i) = x(n) * h(n)$$



Programmability:

1. Native Python programming
2. Native control flow representation

Compilation layer



Performance:

1. Compilation optimization
2. Operator convergence

Operator layer



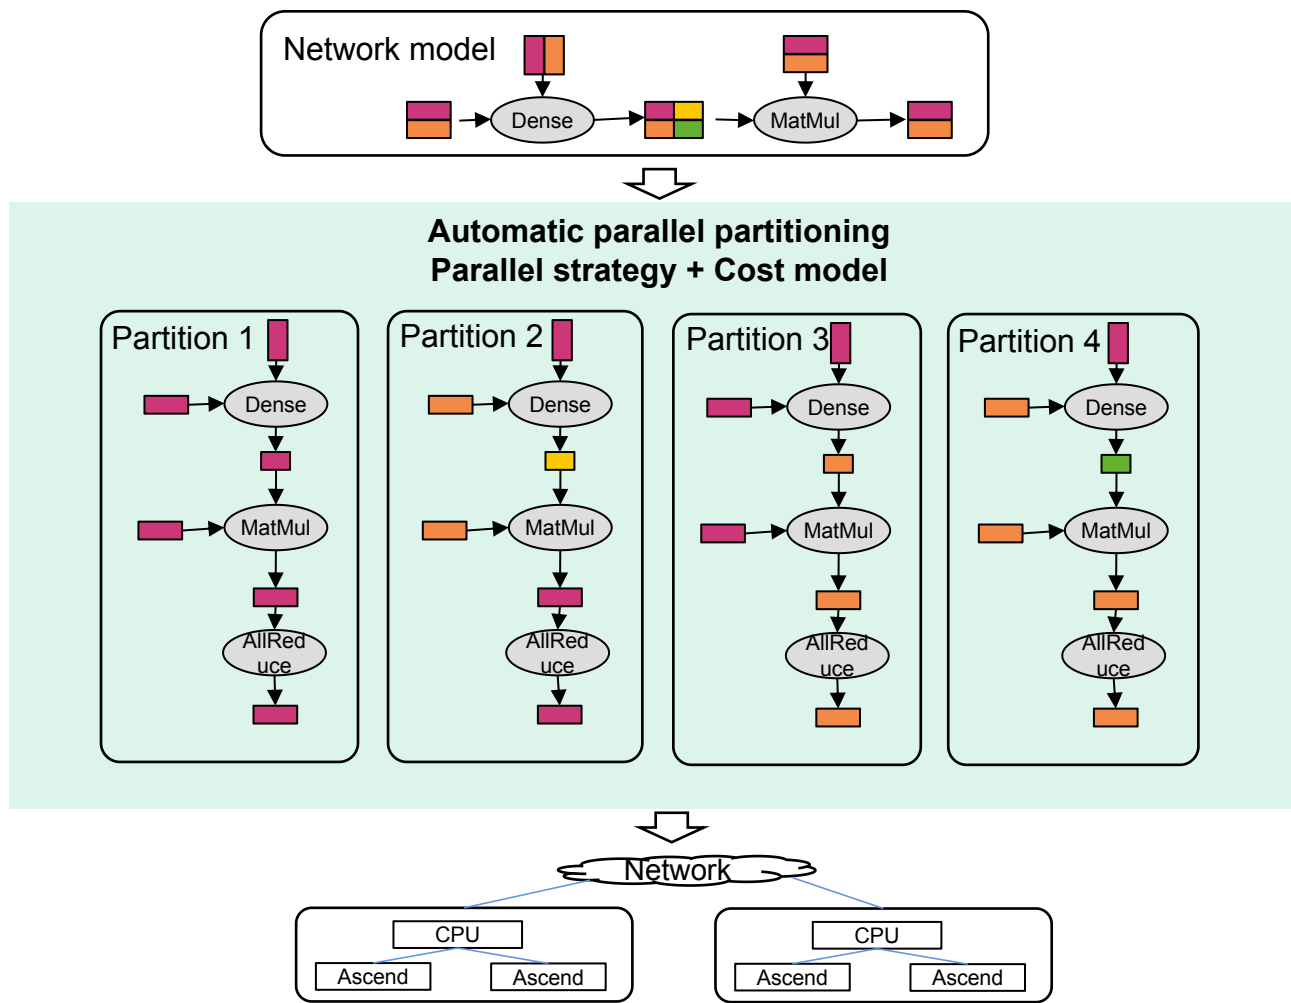
Debugging:

1. Visualization
2. Dynamic execution

Auto Parallel

Challenges:

- ❖ The sizes of datasets and training models keep increasing. Data parallel training reaches the bottleneck due to limited memory of a single device, and model parallel processing is required.
- ❖ Manual Model Parallelism – Difficult to Partition.



1. Automatic operator partitioning
2. Automatic graph partitioning
3. Network topology aware scheduling
4. Automatic search for optimal parallel strategy

Ascend Native Graph Execution


Ascend Native's execution engine

Main challenges in execution



The complexity of AI computing and the diversity of computing power

- ① CPU core, CUBE unit、vector computation
- ② Scalar, vector and tensor computation
- ③ Mixed precision computation
- ④ Dense and sparse matrix computation

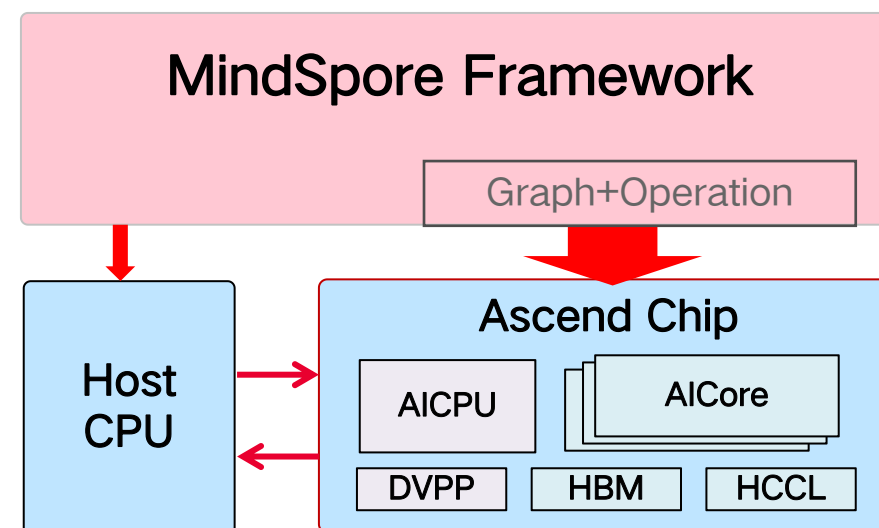


Multi-card operation: Large overhead for parallel control

It is difficult to increase the performance linearly with the number of nodes.

On-Device Implementation

Unload the full computation graph to take advantage of the Ascend chips' s full potential computing power



Visualization



MindSpore

```
from mindspore.ops import functional as F

class CrossEntropyLoss(nn.Cell):
    """
    Define loss for network
    """
    def __init__(self):
        super(CrossEntropyLoss, self).__init__()
        self.cross_entropy = P.SoftmaxCrossEntropyWithLogits()
        self.mean = P.ReduceMean()
        self.one_hot = P.OneHot()
        self.on_value = Tensor(1.0, mstype.float32)
        self.off_value = Tensor(0.0, mstype.float32)
        self.scalar_summary = P.ScalarSummary()

    def construct(self, logits, label):
        label = self.one_hot(label, F.shape(logits)[1], self.on_value, self.off_value)
        loss = self.cross_entropy(logits, label)[0]
        loss = self.mean(loss, (-1,))
        self.scalar_summary("loss", loss)
        return loss
```

```
class AlexNet(nn.Cell):
    """
    Alexnet
    """
    def __init__(self, num_classes=1001):
        super(AlexNet, self).__init__()
        self.batch_size = 32
        self.conv1 = conv(3, 96, 11, stride=4)
        self.conv2 = conv(96, 256, 5, pad_mode="same")
        self.conv3 = conv(256, 384, 3, pad_mode="same")
        self.conv4 = conv(384, 384, 3, pad_mode="same")
        self.conv5 = conv(384, 96, 3, pad_mode="same")
        self.relu = nn.ReLU()
        # self.max_pool2d = nn.MaxPool2d(kernel_size=3, stride=2)
        self.max_pool2d = P.MaxPool(ksize=3, strides=2)
        self.reshape = P.Reshape()
        self.fc1 = fc_with_initialize(6*6*96, 4096)
        self.fc2 = fc_with_initialize(4096, 4096)
        self.fc3 = fc_with_initialize(4096, num_classes)
        self.image_summary = P.ImageSummary()


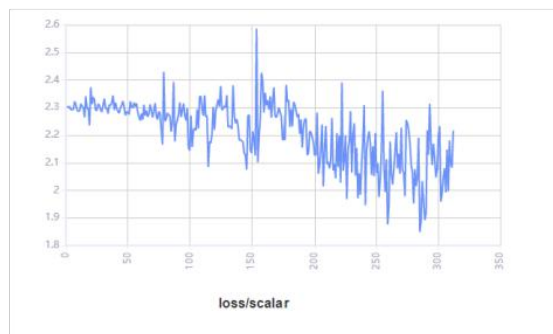
    def construct(self, x):
        self.image_summary("image", x)
        x = self.conv1(x)
        x = self.relu(x)
```

```
class Momentum(Optimizer):
    """
    """
    # Initialize ScalarSummary.
    self.sm_scalar = P.ScalarSummary()

    # prepare summary operator
    self.histogram_summary = P.HistogramSummary()
    # prepare tags
    self.weight_names = [param.name for param in self.parameters]
    self.gradient_names = [param.name + ".gradient" for param in self.parameters]
    self.param_count = len(self.parameters)

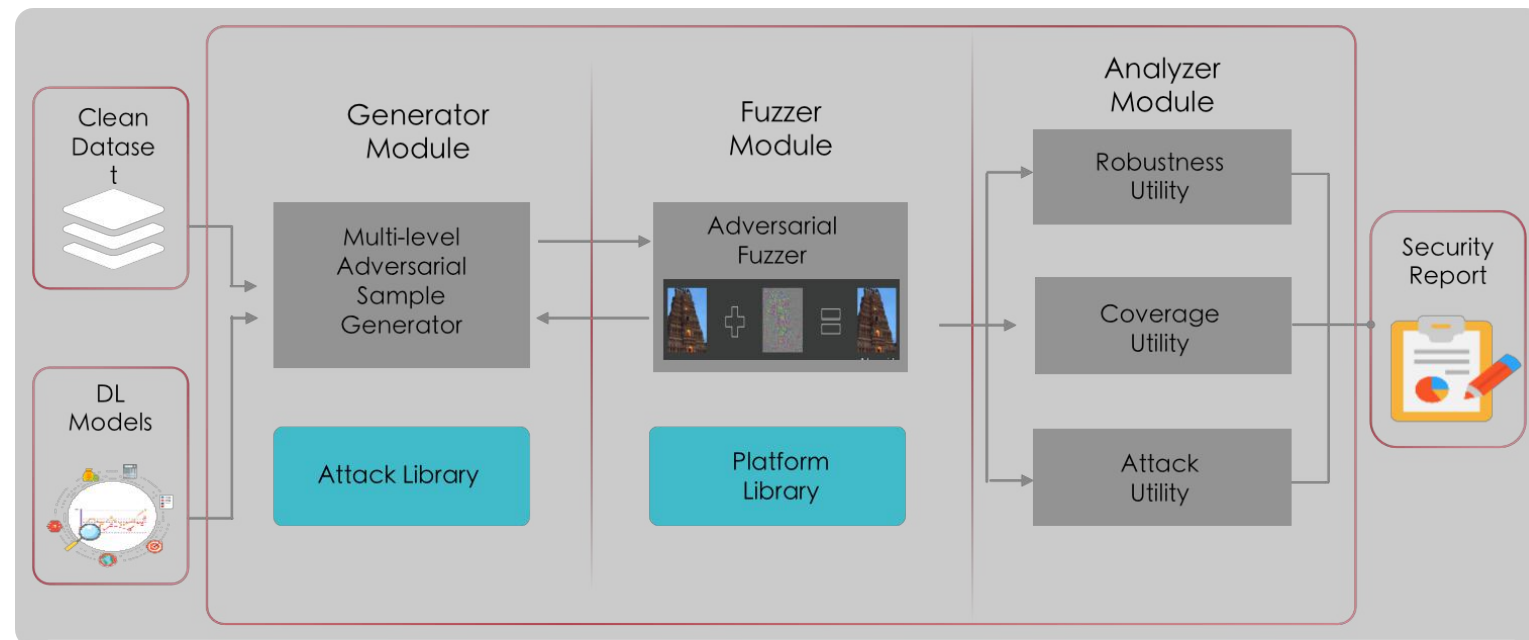
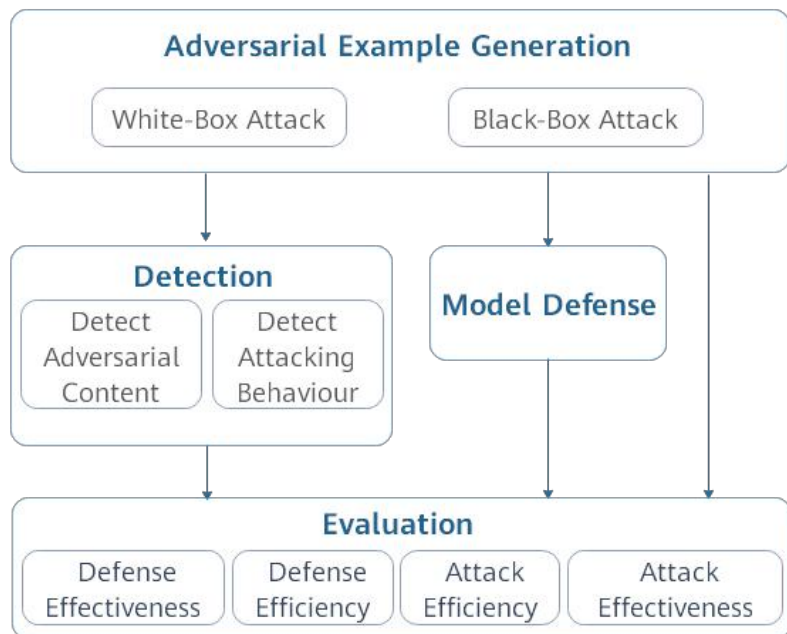
    def construct(self, grads):
        # record weights
        for i in range(self.param_count):
            self.histogram_summary(self.weight_names[i], self.params[i])

        # record gradients
        for i in range(self.param_count):
            self.histogram_summary(self.gradient_names[i], grads[i])
```



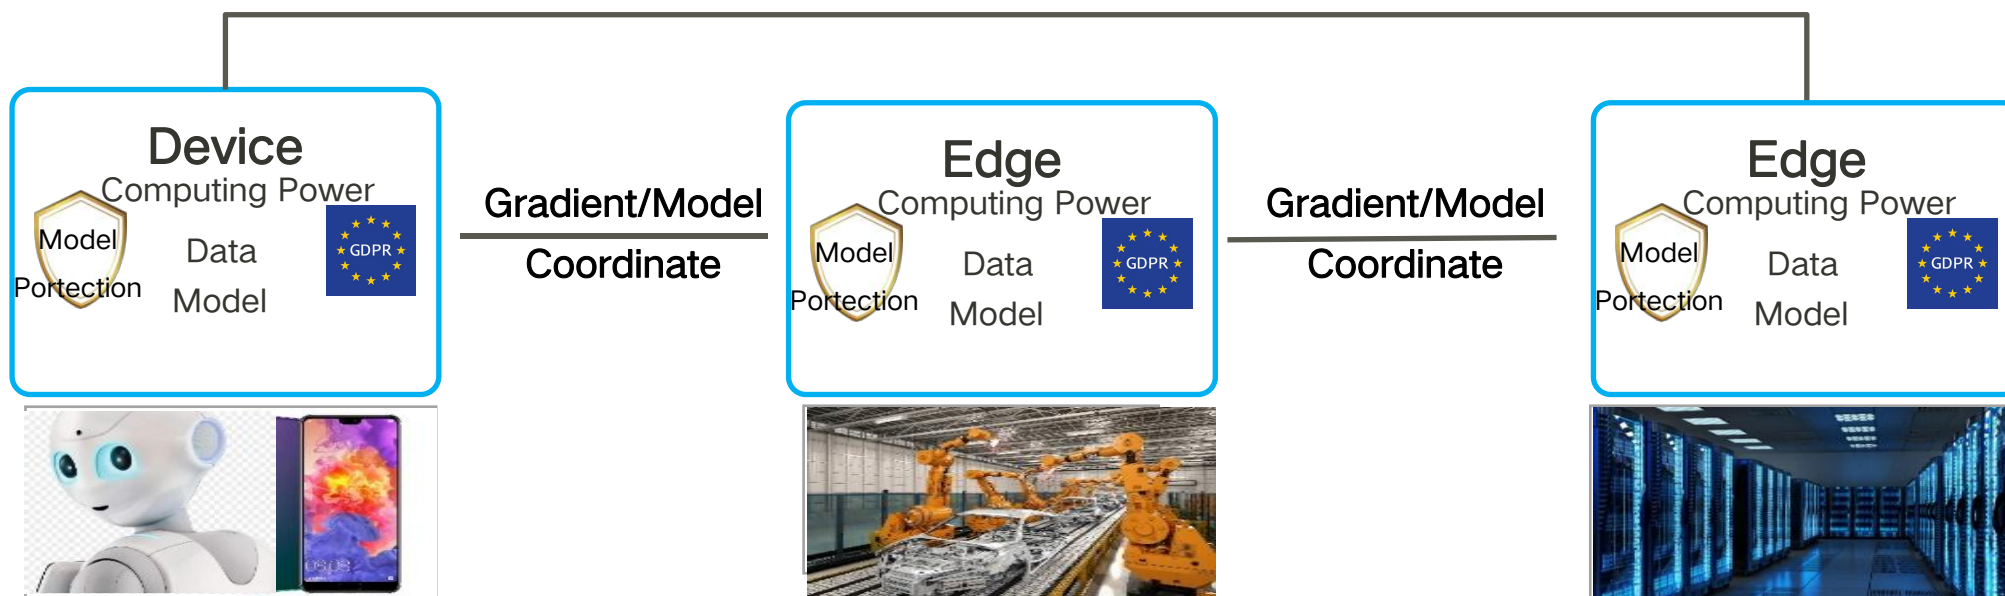
Robustness Evaluation

MindArmour

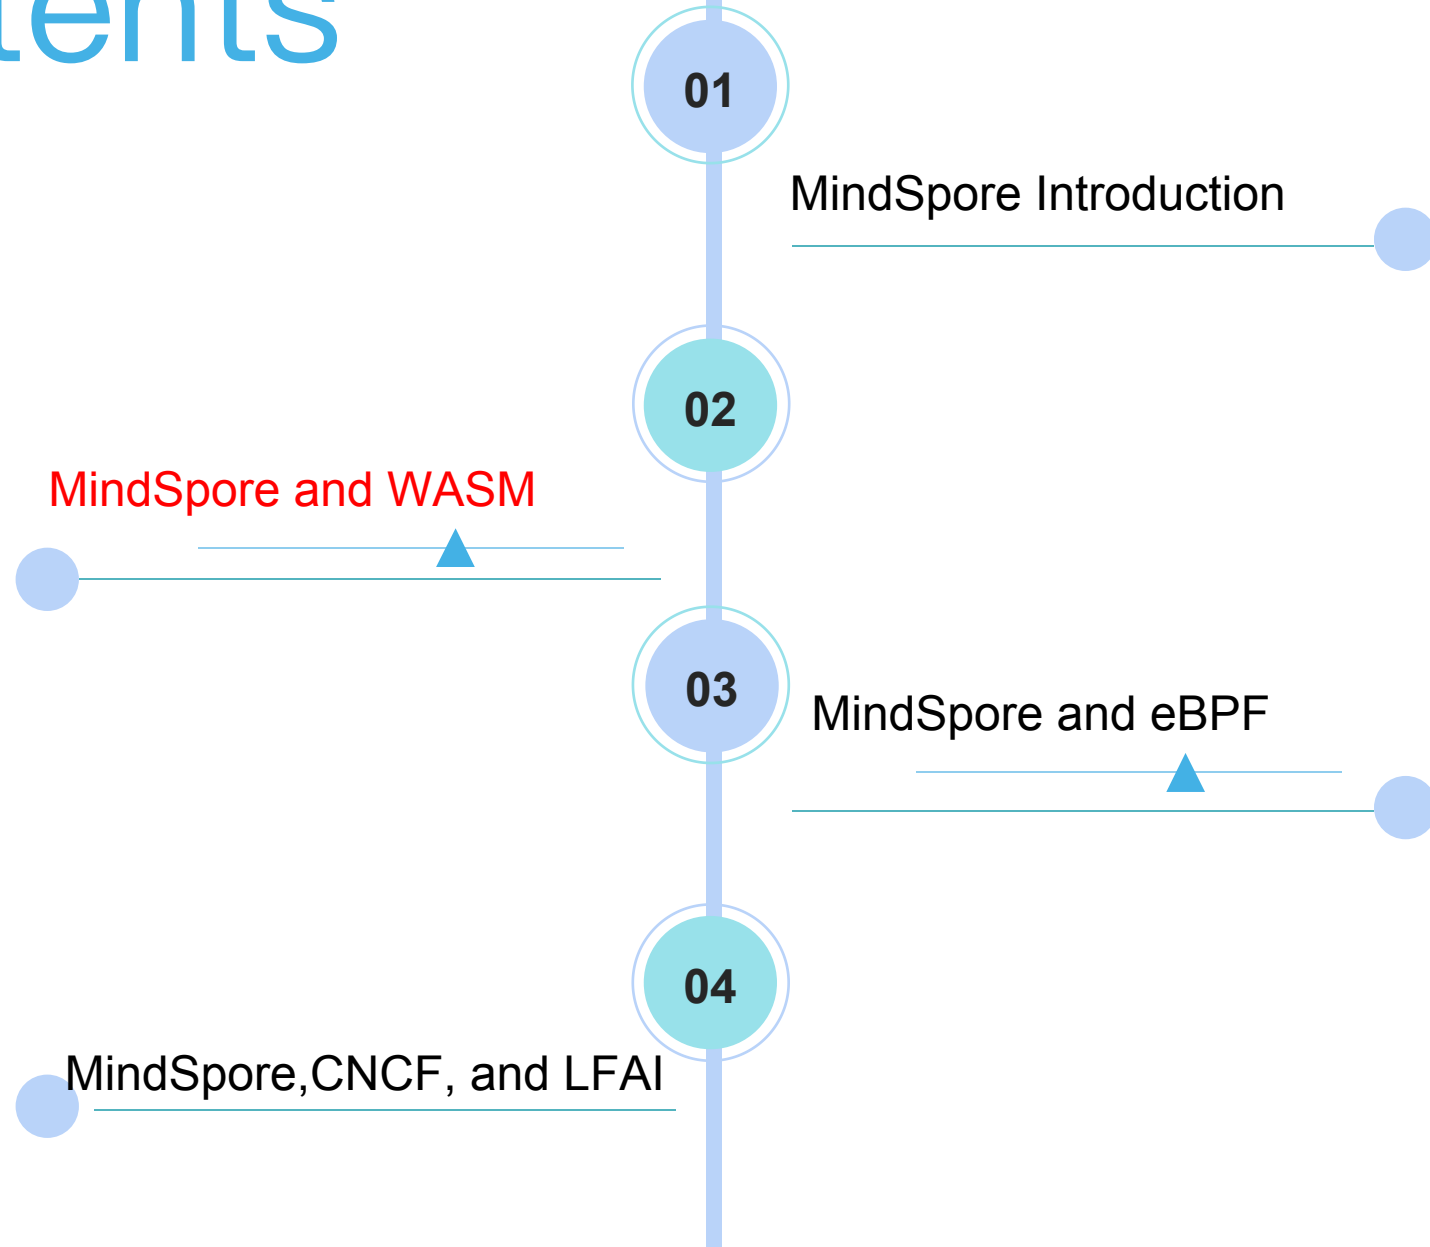


All Scenario Collaboration

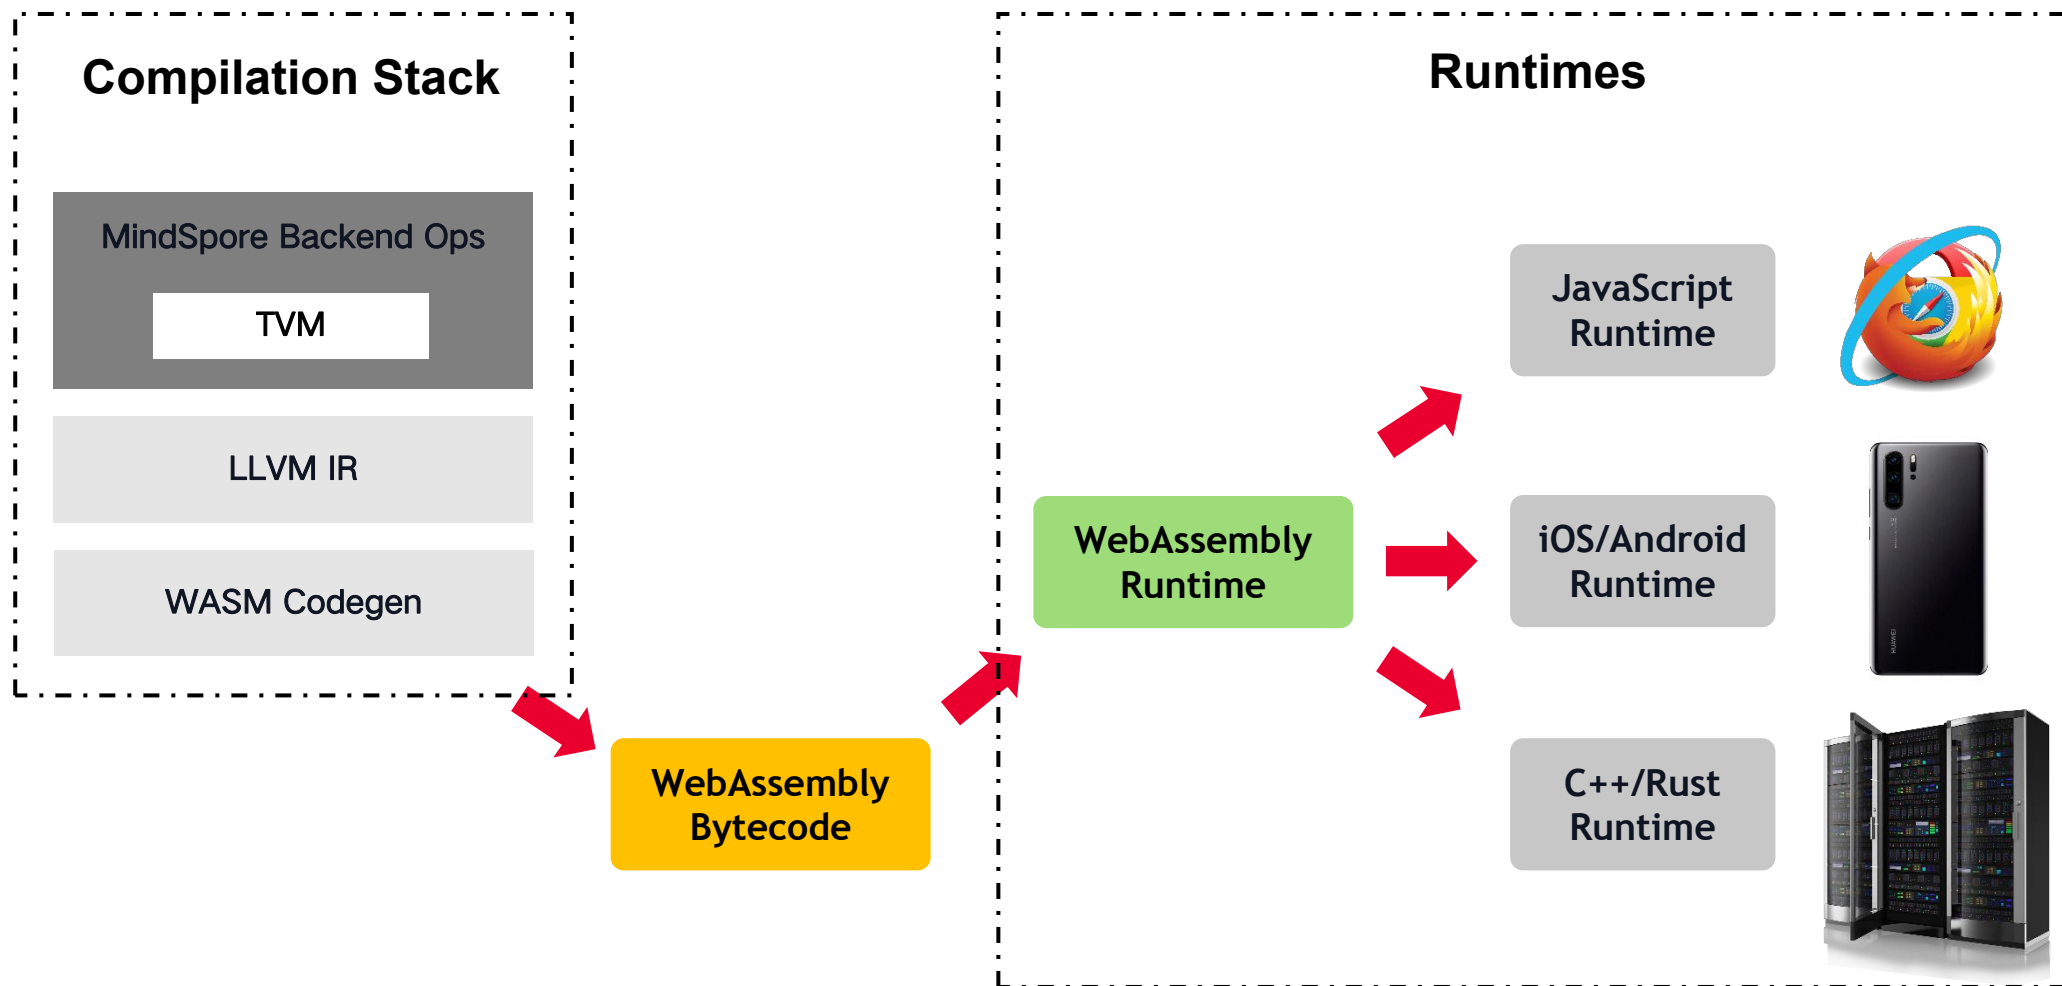
End-to-edge-cloud scenarios are built on demand, unified AI architecture, size flexible to various enterprise environment, develop once and deploy anywhere



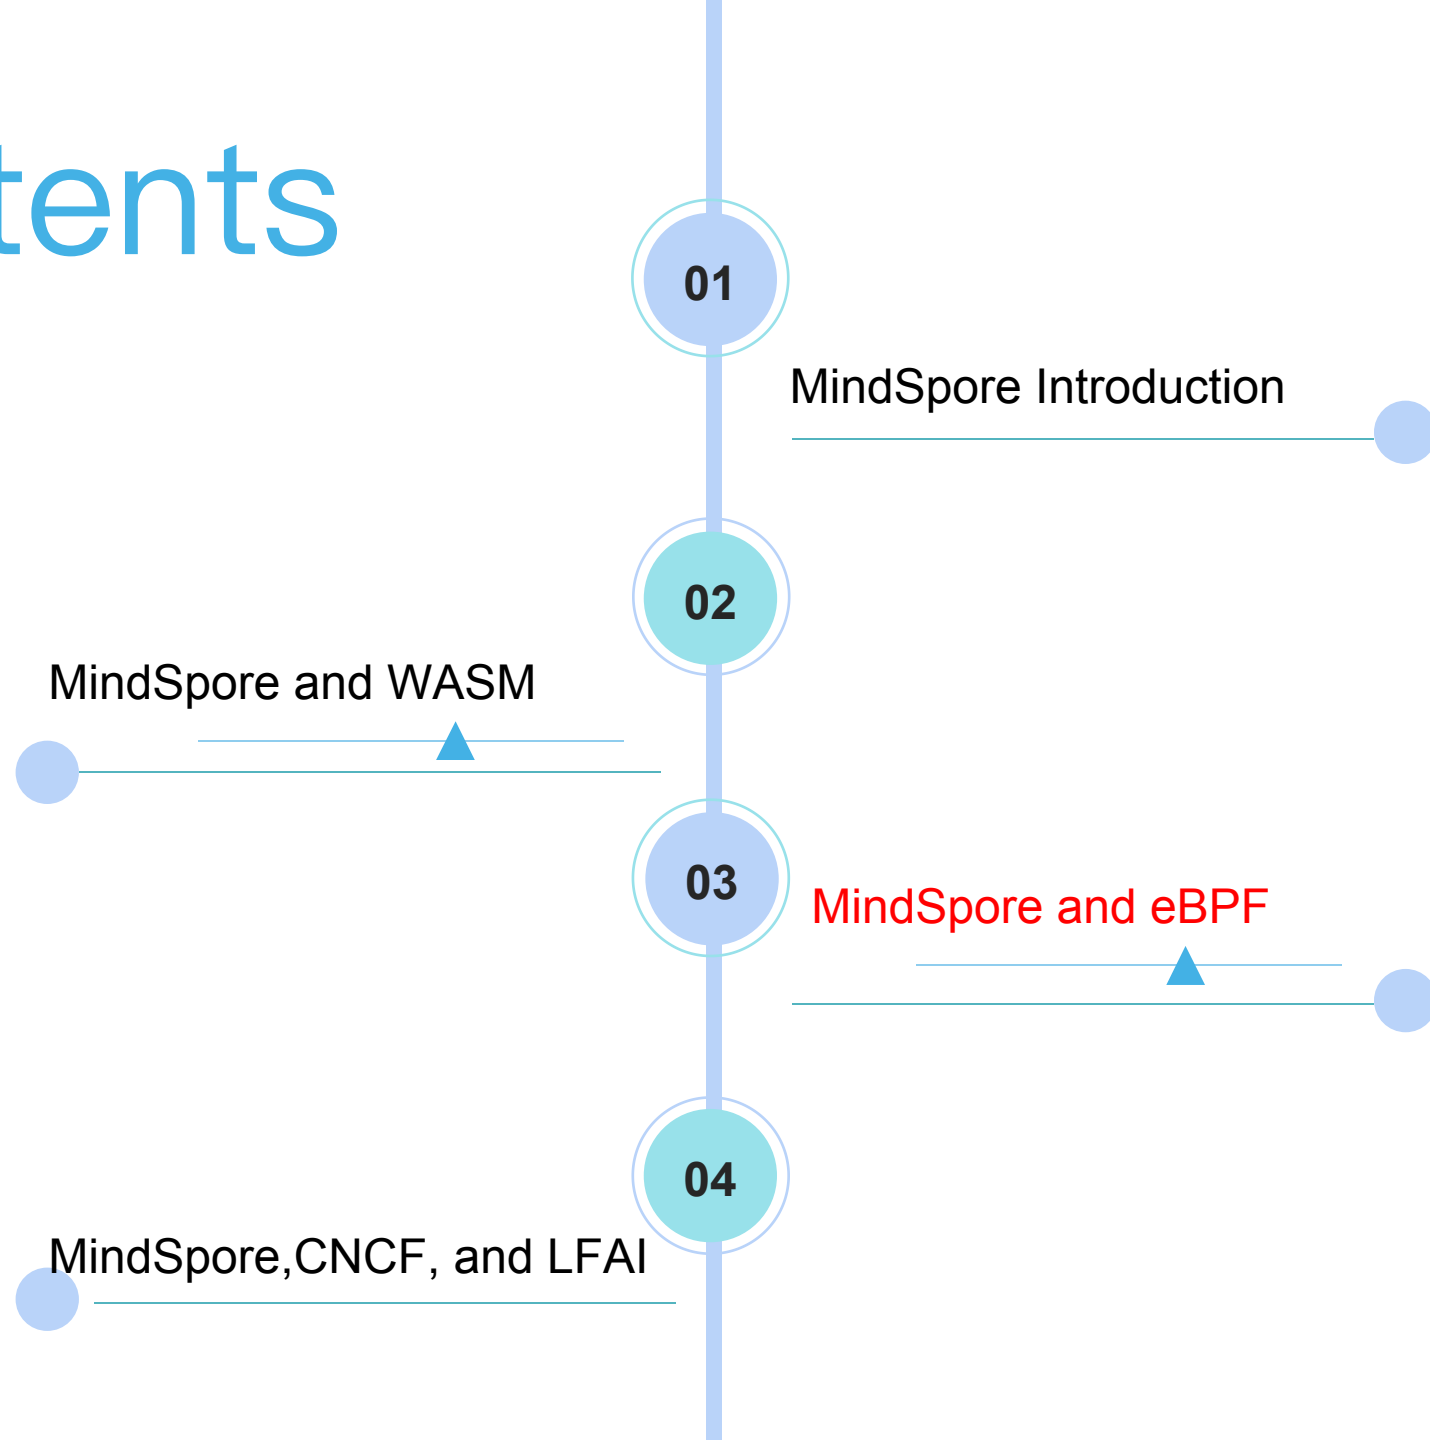
Contents




MindSpore and WASM



Contents



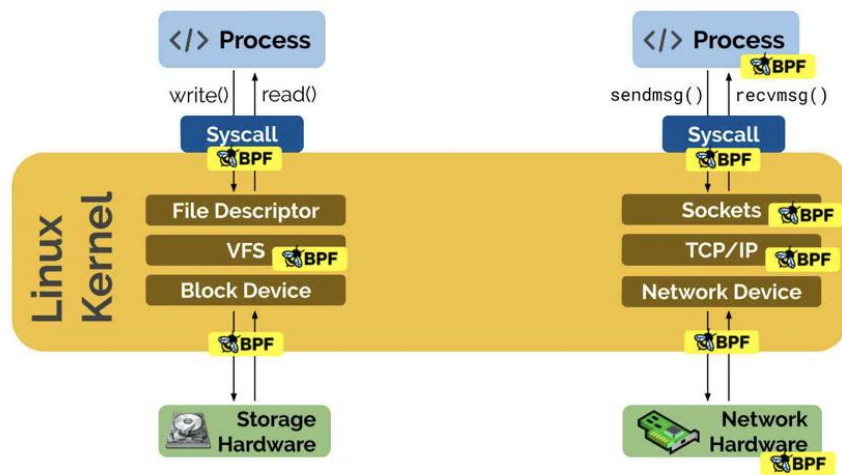


MindSpore

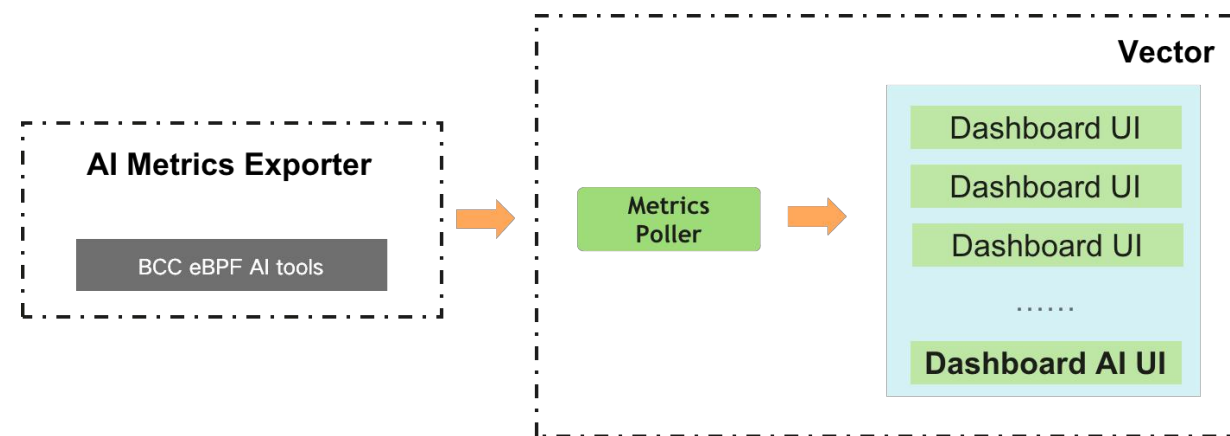
MindSpore and eBPF



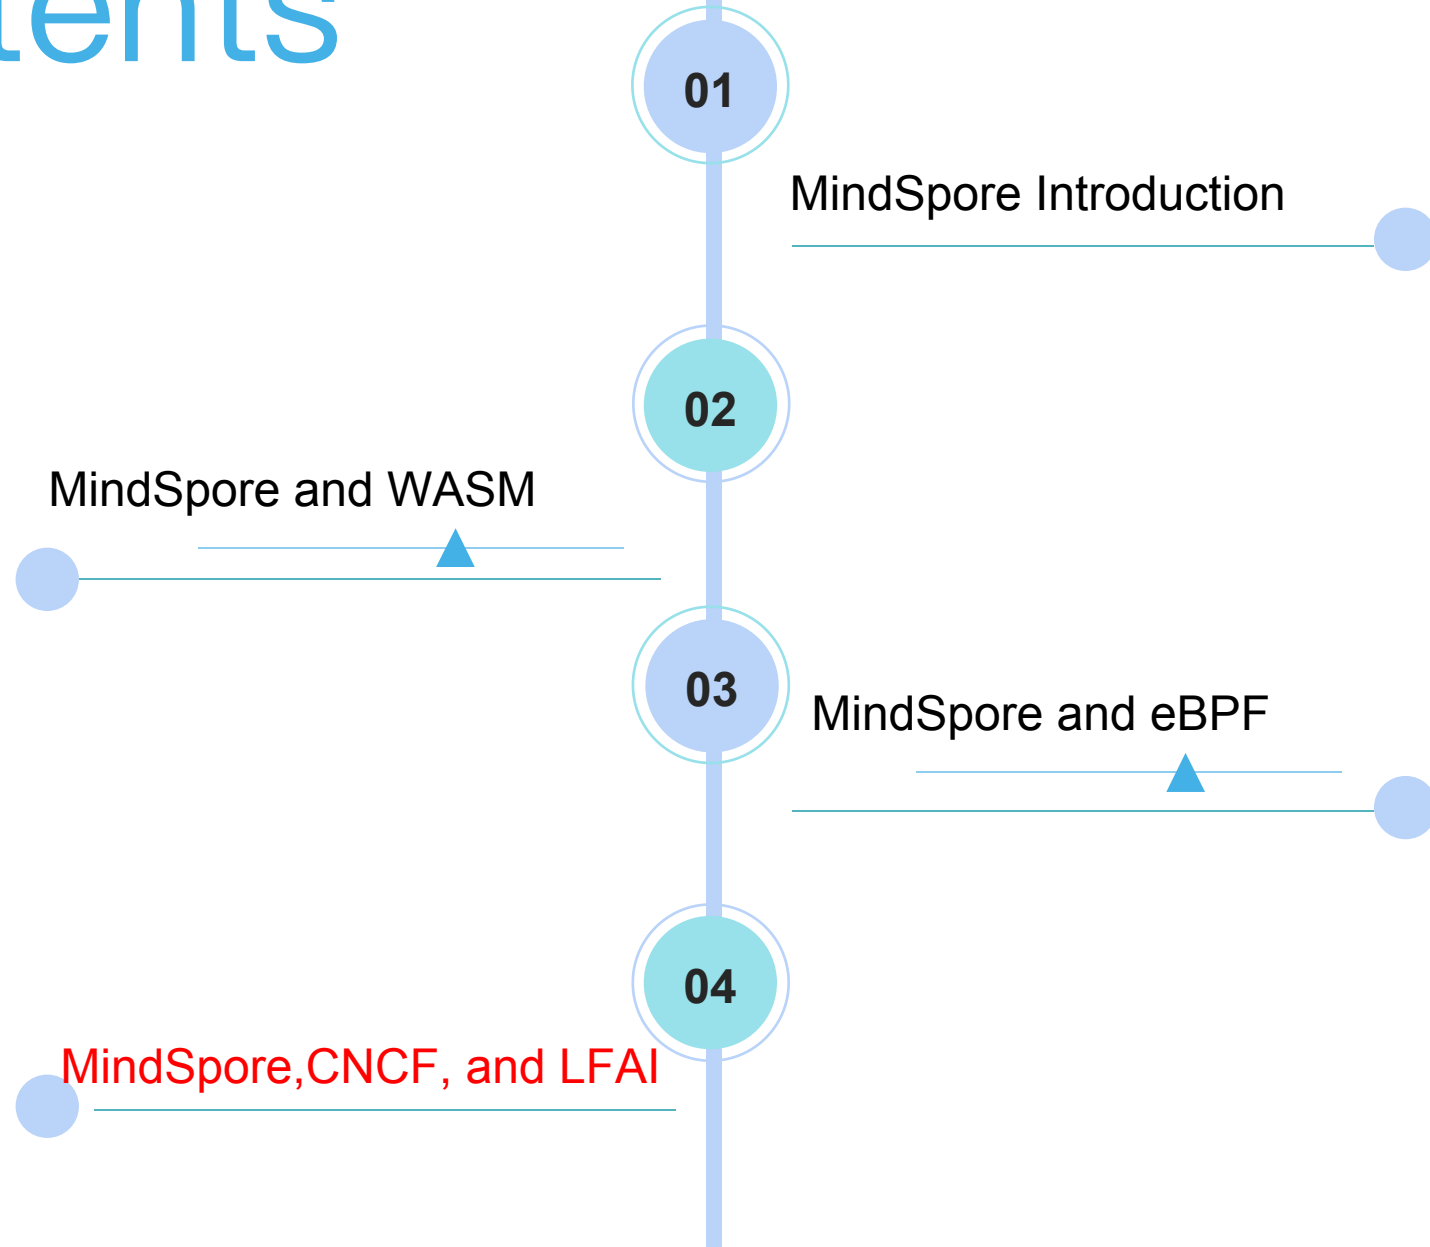
eBPF Hooks



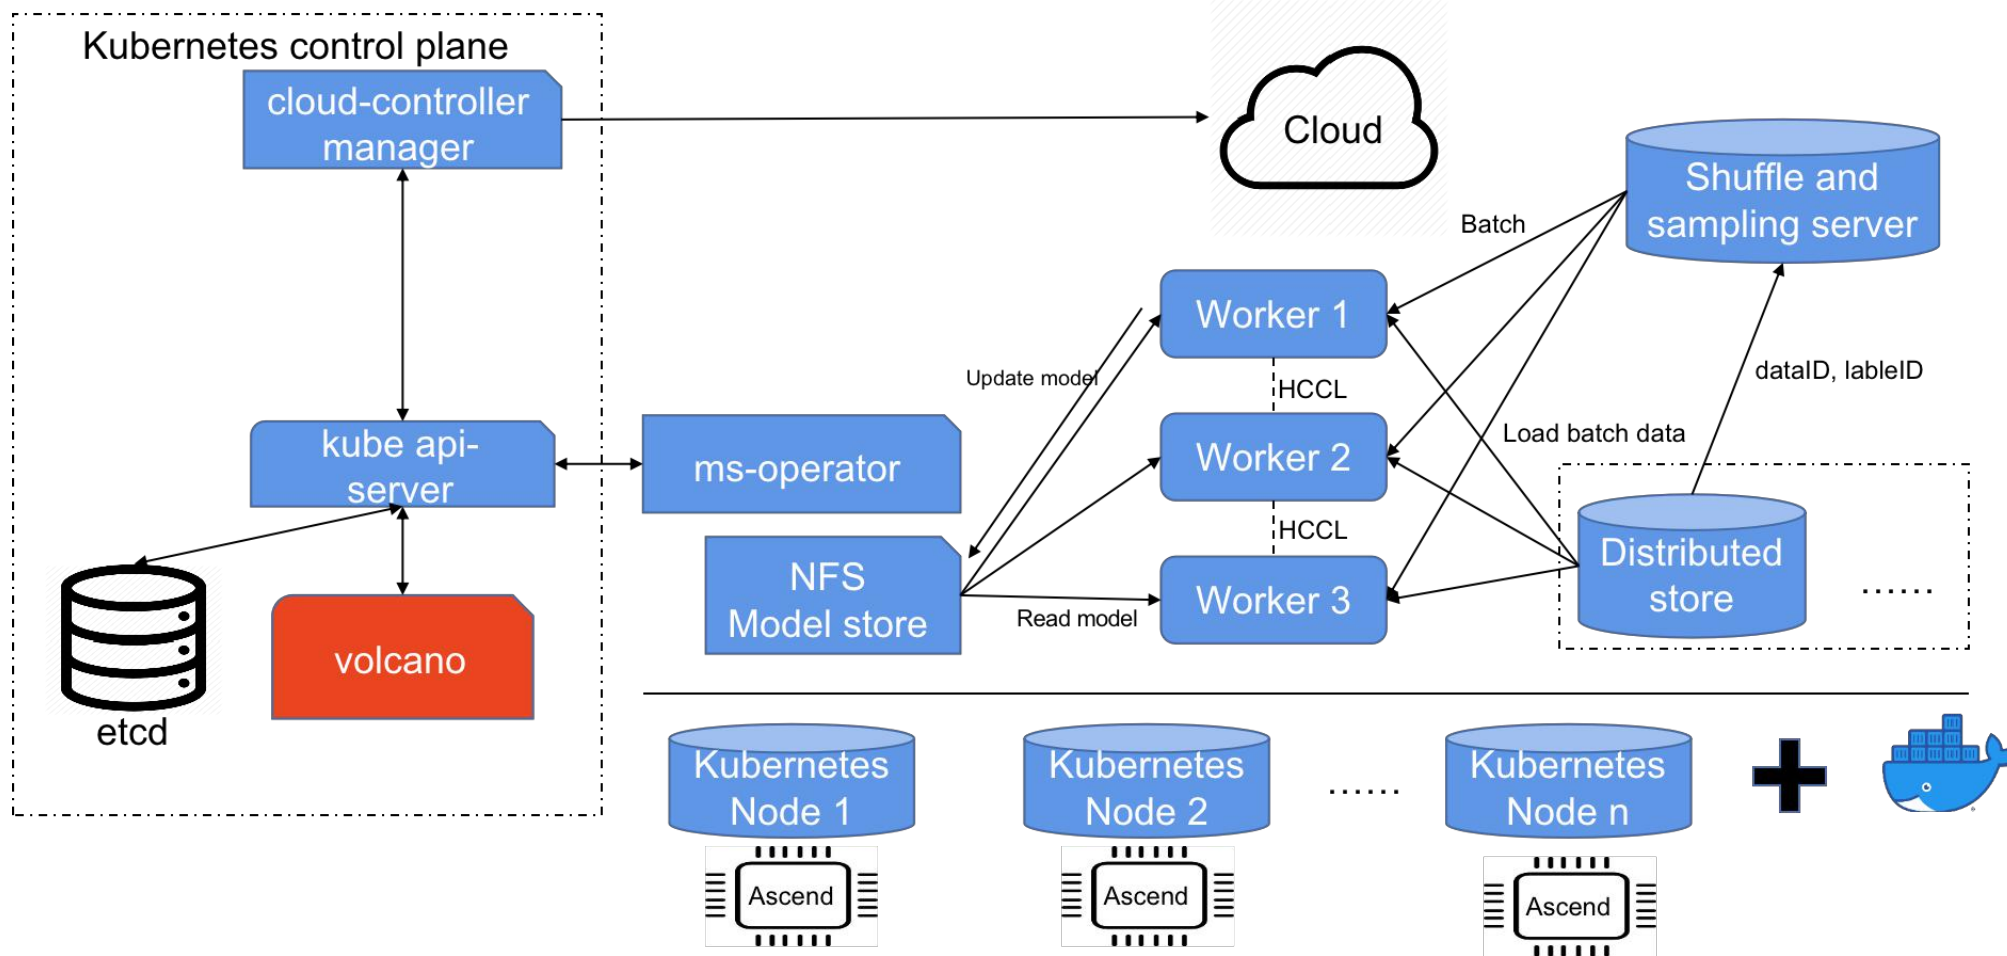
Where can you hook? kernel functions (kprobes), userspace functions (uprobes), system calls, fentry/fexit, tracepoints, network devices (tc/xdp), network routes, TCP congestion algorithms, sockets (data level)



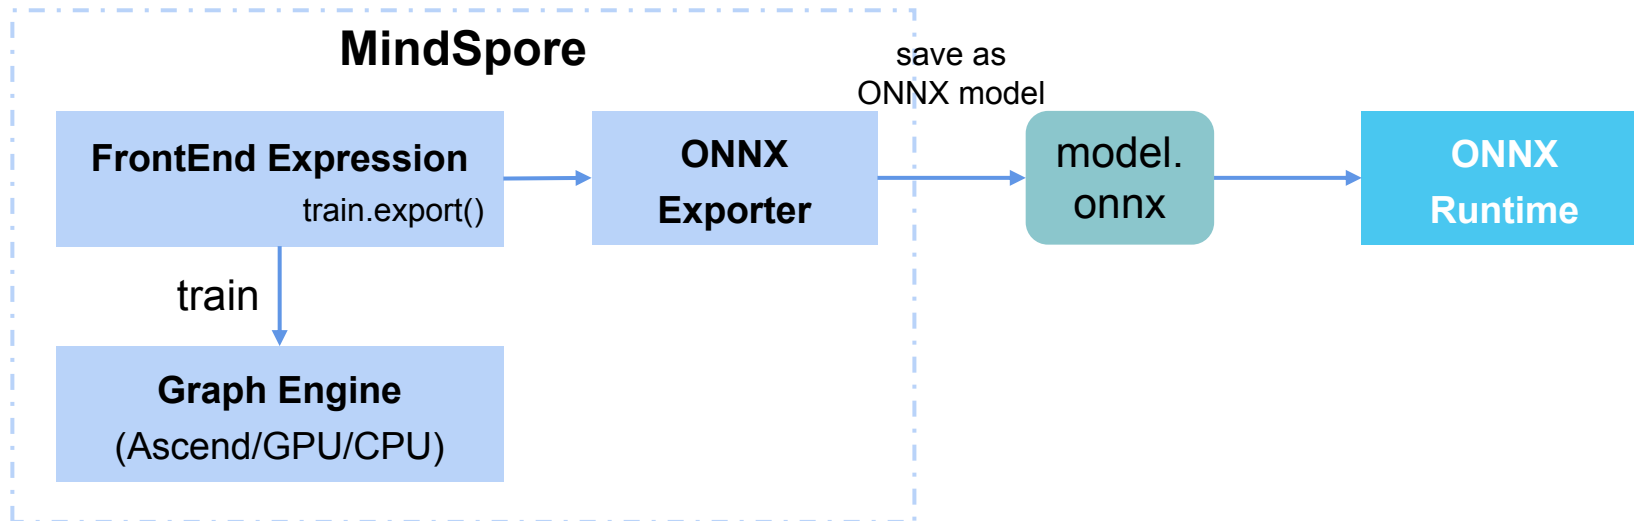
Contents



MindSpore and Kubeflow



MindSpore ONNX Exporter Introduction



1. Use MindSpore model train API to perform model training with saving checkpoint parameters
2. Load model parameters into the network to be exported (such like LeNet)
3. Call `train.export()` to convert MindSpore model to ONNX model
4. Perform model inference on ONNX Runtime



MindSpore

MindSpore ONNX Exporter Support

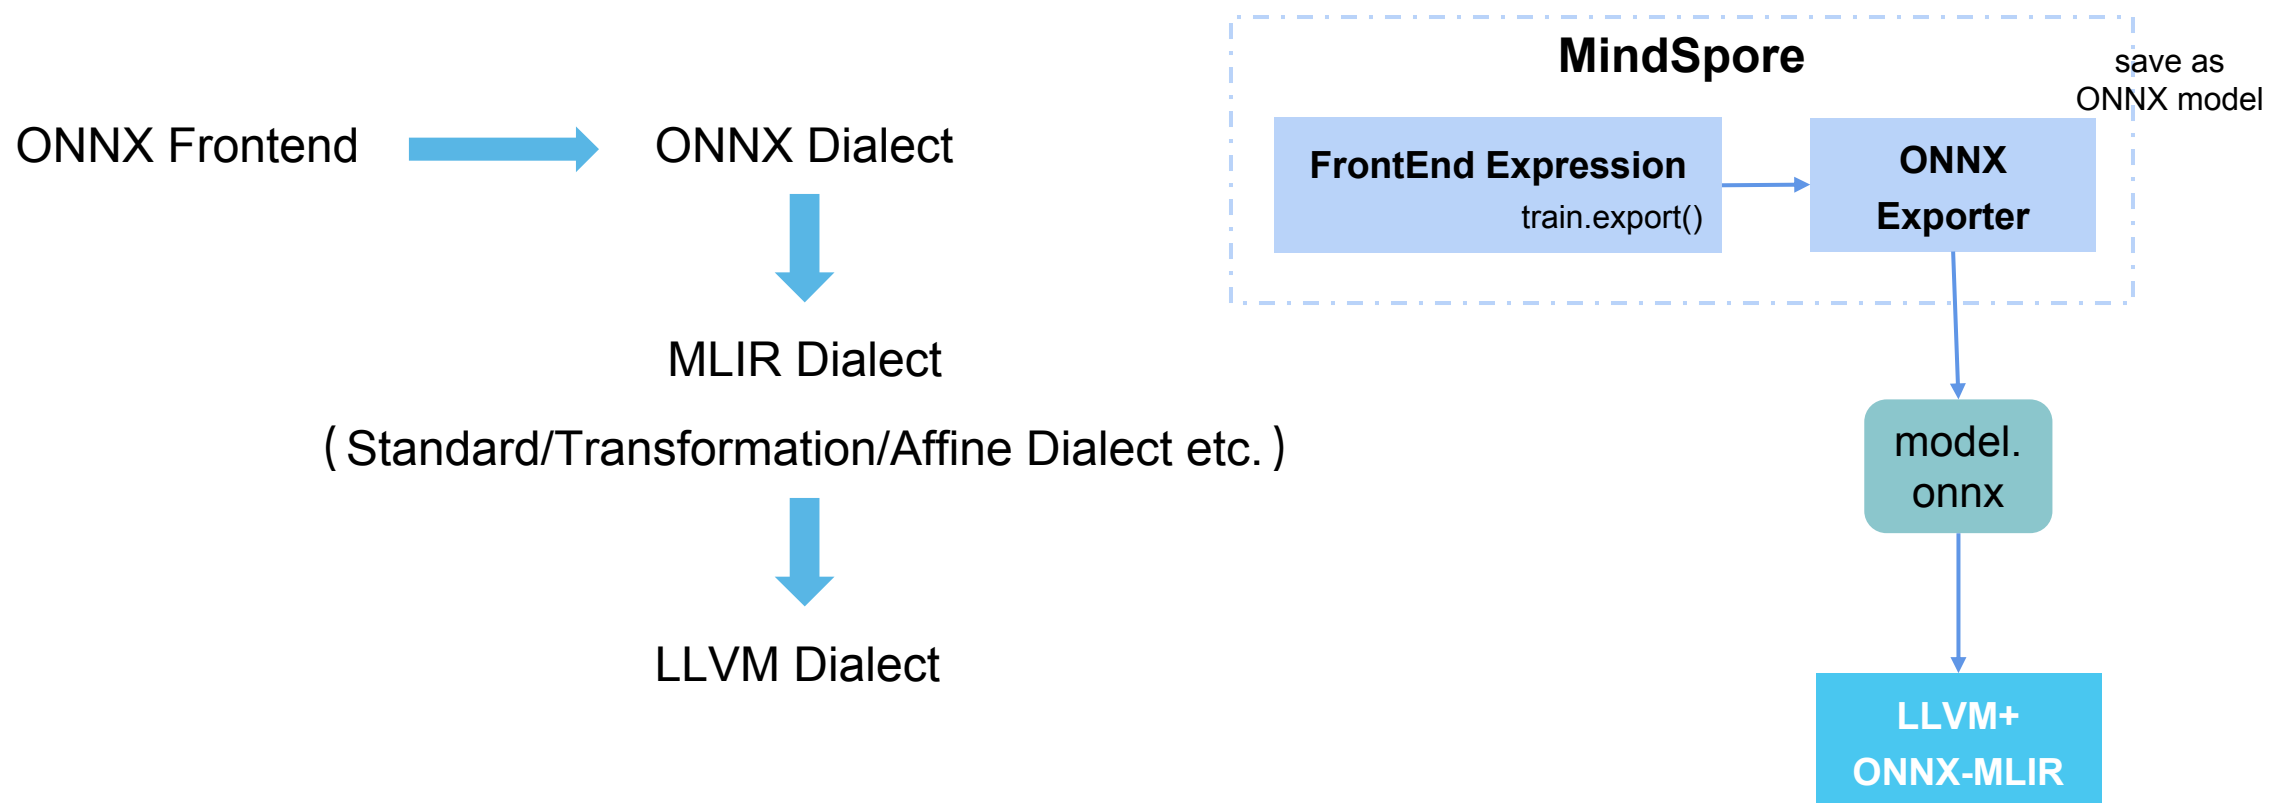
MindSpore Operator	ONNX Operator	MindSpore Operator	ONNX Operator
TensorAdd	Add	BatchNorm	BatchNormalization
Mul	Mul	Reshape	Reshape
ReLU	Relu	ReduceMean	ReduceMean
Sigmoid	Sigmoid	Cast	Cast
Flatten	Flatten	PReLU	PRelu
Squeeze	Squeeze	Argmax	ArgMax
Conv2D	Conv	SimpleMean	AveragePool
BiasAdd	Add	MaxPool	MaxPool
MatMul	Gemm	AvgPool	AveragePool

- **Operator Support:** 18
- **Network Support:** LeNet、ResNet50、AlexNet
- **Roadmap:** More operators and networks (in ModelZoo) are WIP



MindSpore

Exploring MindSpore and ONNX-MLIR (Idea)




<https://github.com/onnx/onnx-mlir>


MindSpore and LFAI MLworkflow & Interop Committee



- 1 Northbound Interoperability when the AI Native programming framework is adopted for applications in different areas.
- 2 Southbound Interoperability when the AI Native programming framework is used on various compute and storage backends
- 3 Framework Interoperability when traditional deep learning frameworks need to be supported or transformed. (**ONNX**)
- 4 Build Interoperability when various types of deployment could be reproducible on different pipelines. (**Kubeflow**)



MindSpore and LFAI MLworkflow & Interop Committee



Northbound: Interop Among Frameworks (IR, Crypto, Type System, Metadata ...)



AI Native Programming
(Pyro, Julia, MindSpore, ...)



Southbound: Interop Between Frameworks and Backends



AI Domain Specific Hardware
(Ascend, GPUs, FPGAs, ASICs,...)



Storage
(Raw Volumes, Object Storage, Database...)



Call For Participation



- Checkout the Code
 - ❑ <https://gitee.com/mindspore> (main development)
 - ❑ <https://github.com/mindspore-ai> (mirror)
- Try with docker
 - ❑ `docker pull mindspore/mindspore-cpu:0.5.0-beta`
 - ❑ `docker pull mindspore/mindspore-gpu:0.5.0-beta`
- Discussion
 - ❑ Slack: send email to zhipengh512@gmail.com for channel invitation link
 - ❑ Mailing list: <https://mailweb.mindspore.cn/postorius/lists/mindspore-discuss.mindspore.cn/>

Thank You



MindSpore

<https://www.mindspore.cn>



<https://www.gitee.com/mindspore>



<https://www.github.com/mindspore-ai>

LF AI General Updates




Project Updates

Projects













Graduation

Incubation

Graduated LFAI Projects (3)

 Acumos LF Artificial Intelligence Foundation ★10	 Angel Angel-ML LF Artificial Intelligence Foundation ★5,896	 ONNX LF Artificial Intelligence Foundation ★8,693
-------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------

Incubating LFAI Projects (12)

 ADLIK Adlik LF Artificial Intelligence Foundation ★115	 IBM Adversarial Robustness Toolkit (ART) Adversarial Robustness Toolkit LF Artificial Intelligence Foundation ★1,597	 IBM AI Explainability 360 (AIX360) AI Explainability 360 Toolkit LF Artificial Intelligence Foundation ★586	 IBM AI Fairness 360 Toolkit (AIF360) AI Fairness 360 Toolkit LF Artificial Intelligence Foundation ★996
 EDL Elastic Deep Learning Elastic Deep Learning (EDL) LF Artificial Intelligence Foundation ★88	 ForestFlow ForestFlow LF Artificial Intelligence Foundation ★32	 HOROVOD Horovod LF Artificial Intelligence Foundation ★9,579	 MARQUEZ Marquez LF Artificial Intelligence Foundation ★319
 Milvus Milvus LF Artificial Intelligence Foundation ★3,698	 NNStreamer NNStreamer LF Artificial Intelligence Foundation ★246	 Pyro Pyro LF Artificial Intelligence Foundation ★6,316	 sparklyr sparklyr LF Artificial Intelligence Foundation ★745

Companies hosting projects in LF AI



FACEBOOK



Looking to host a project with LF AI

Hosted project stages and life cycle:

<https://lfai.foundation/project-stages-and-lifecycle/>

Offered services for hosted projects:

<https://lfai.foundation/services-for-projects/>

Contact:

Jim Spohrer (TAC Chair) and Ibrahim Haddad (ED, LF AI)

Promoting Upcoming Project Releases

We promote project releases via a blog post and on LF AI [Twitter](#) and/or [LinkedIn](#) social channels

For links to details on upcoming releases for LF AI hosted projects visit the [Technical Project Releases wiki](#)

If you are an LF AI hosted project and would like LF AI to promote your release, reach out to pr@lfai.foundation to coordinate in advance (min 2 wks) of your expected release date.

Outreach Committee

LF AI PR/Comms

- › Please follow LF AI on [Twitter](#) & [LinkedIn](#) and help amplify news via your social networks - Please retweet and share!
 - › Also watch for news updates via the tac-general mail list
 - › View recent announcement on the [LF AI Blog](#)
- › Open call to publish project/committee updates or other relevant content on the [LF AI Blog](#)
- › To discuss more details on participation or upcoming announcements, please email pr@lfai.foundation

Events

- › Upcoming Events
 - › Visit the [LF AI Events Calendar](#) or the [LF AI 2020 Events wiki](#) for a list of all events
 - › To participate visit the [LF AI 2020 Events wiki page](#) or email info@lfai.foundation
- › Please consider holding virtual events
 - › To discuss participation, please email events@lfai.foundation

Upcoming Events



OSOW
OPEN SOURCE CHINA
OPEN SOURCE WORLD

OPU 中国开源软件(OSS)推进联盟
China OSS Promotion Union

**The 15th Open Source China
Open Source World Summit**
第十五届“开源中国开源世界”高峰论坛

研发基于开源深度信息技术
重构新业态 拓展新生态

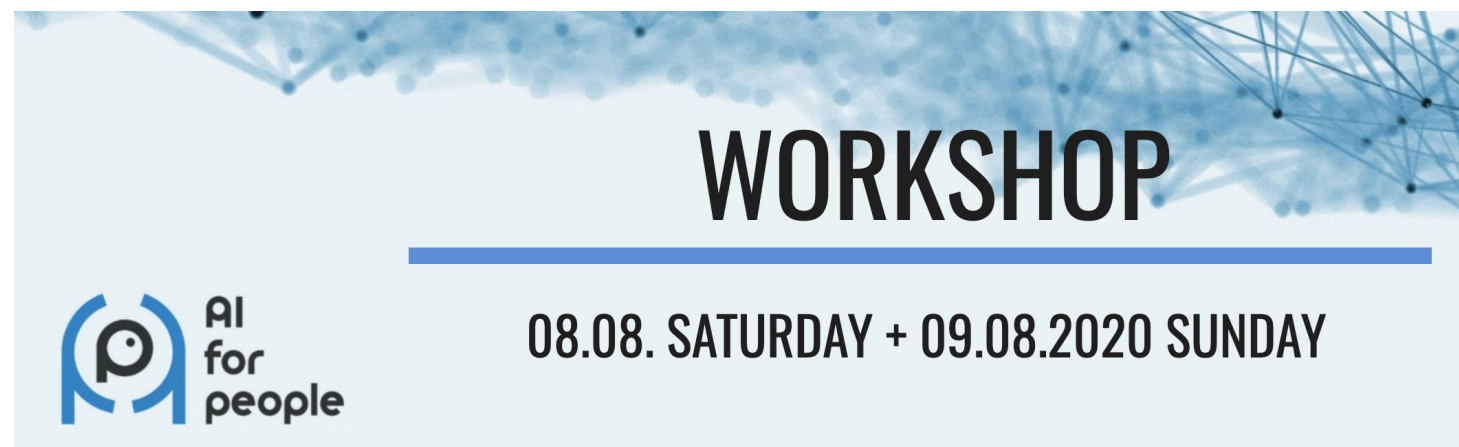
2020年7月15日线上直播






CLOUD NATIVE + OPEN SOURCE

Virtual Summit China 2020



WORKSHOP

 AI
for
people

08.08. SATURDAY + 09.08.2020 SUNDAY

Recent Events

- › **Virtual LF AI Day EU**
 - › June 22
 - › Post event summary [Blog](#)

- › Thank you to Orange for hosting!



Recent Events

- › **LF AI at OSS NA**
 - › June 29 - July 1
 - › Post event summary [Blog](#)
- › Thank you to our speakers and booth volunteers!

Connect with    **OPEN SOURCE SUMMIT**
NORTH AMERICA

Visit us at the LF AI Booth
Bronze Hall
June 29-July 1

Join us at the LF AI Mini Summit
Thursday, July 2
10:00-11:30 AM CDT

**Attend sessions in the
AI/ML/DL Track**
June 29-July 1

events.linuxfoundation.org/open-source-summit-north-america/register/

Call to Participate in Ongoing Efforts

Trusted AI

- › **Leadership:**
Animesh Singh (IBM), Souad Ouali (Orange), and Jeff Cao (Tencent)
- › **Goal:** Create policies, guidelines, tooling and use cases by industry
- › **Github:**
<https://github.com/lfai/trusted-ai>
- › **Wiki:**
<https://wiki.lfai.foundation/display/DL/Trusted+AI+Committee>
- › **To participate:**
<https://lists.lfai.foundation/g/trustedai-committee/>
- › **Next call:** Bi-weekly on Thursdays at 7am PT, subscribe to group calendar on wiki
<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091895>

ML Workflow & Interop

- › **Leadership:**
Huang “Howard” Zhipeng (Huawei)
- › **Goal:**
Define an ML Workflow and promote cross project integration
- › **Wiki:**
<https://wiki.lfai.foundation/display/DL/ML+Workflow+Committee>
- › **To participate:**
<https://lists.lfai.foundation/g/mlworkflow-committee>
- › **Next call:** Every 4 weeks on Thursdays at 7:00 am PT, subscribe to group calendar on wiki
<https://wiki.lfai.foundation/pages/viewpage.action?pageId=18481242>

Launching an effort to create AI Ethics Training

Initial developed course by the LF: Ethics in AI and Big Data - published on edX platform:
<https://www.edx.org/course/ethics-in-ai-and-big-data>

The goal is to build 2 more modules and package all 3 as a professional certificate - a requirement for edX

- › The LF would cover the cost of the production and promotion
- › The course would be offered for free
- › The credit of the course will go to content creator and their organizations
- › Initial interested parties: IBM, AI for People, Montreal AI Ethics Institute, Ethical ML Institute
- › **To participate:**
<https://lists.lfai.foundation/g/aiethics-training>

Upcoming TAC Meetings

Upcoming TAC Meetings

- › **July 30:** Guest Presentations
 - › [OpenPower](#) - James Kulina (Executive Director)
 - › [ModelDB](#) - Conrado Silva Miranda (Verta.ai)

- › **August 13:** To be announced

Please send agenda topic requests to tac-general@lists.lfai.foundation

TAC Meeting Details

- › To subscribe to the TAC Group Calendar, visit the wiki: <https://wiki.lfai.foundation/x/XQB2>
- › Join from PC, Mac, Linux, iOS or Android: <https://zoom.us/j/430697670>
- › Or iPhone one-tap:
 - › US: +16465588656,,430697670# or +16699006833,,430697670#
- › Or Telephone:
 - › Dial(for higher quality, dial a number based on your current location):
 - › US: +1 646 558 8656 or +1 669 900 6833 or +1 855 880 1246 (Toll Free) or +1 877 369 0926 (Toll Free)
- › Meeting ID: 430 697 670
- › International numbers available: <https://zoom.us/u/achYtcw7uN>

Open Discussion

Legal Notices

- › The Linux Foundation, The Linux Foundation logos, and other marks that may be used herein are owned by The Linux Foundation or its affiliated entities, and are subject to The Linux Foundation's Trademark Usage Policy at <https://www.linuxfoundation.org/trademark-usage>, as may be modified from time to time.
- › Linux is a registered trademark of Linus Torvalds. Please see the Linux Mark Institute's trademark usage page at <https://lmi.linuxfoundation.org> for details regarding use of this trademark.
- › Some marks that may be used herein are owned by projects operating as separately incorporated entities managed by The Linux Foundation, and have their own trademarks, policies and usage guidelines.
- › TWITTER, TWEET, RETWEET and the Twitter logo are trademarks of Twitter, Inc. or its affiliates.
- › Facebook and the "f" logo are trademarks of Facebook or its affiliates.
- › LinkedIn, the LinkedIn logo, the IN logo and InMail are registered trademarks or trademarks of LinkedIn Corporation and its affiliates in the United States and/or other countries.
- › YouTube and the YouTube icon are trademarks of YouTube or its affiliates.
- › All other trademarks are the property of their respective owners. Use of such marks herein does not represent affiliation with or authorization, sponsorship or approval by such owners unless otherwise expressly specified.
- › The Linux Foundation is subject to other policies, including without limitation its Privacy Policy at <https://www.linuxfoundation.org/privacy> and its Antitrust Policy at <https://www.linuxfoundation.org/antitrust-policy>, each as may be modified from time to time. More information about The Linux Foundation's policies is available at <https://www.linuxfoundation.org>.
- › Please email legal@linuxfoundation.org with any questions about The Linux Foundation's policies or the notices set forth on this slide.