

Hyperledger Fabric 2.0 : What's new?

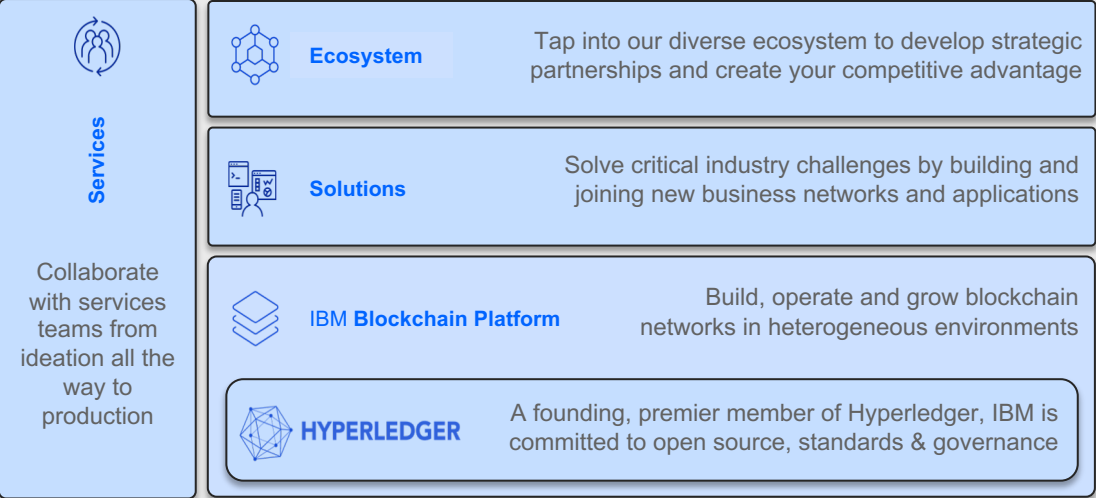
Maciej Jędrzejczyk

CEE Blockchain Technical Leader, IBM

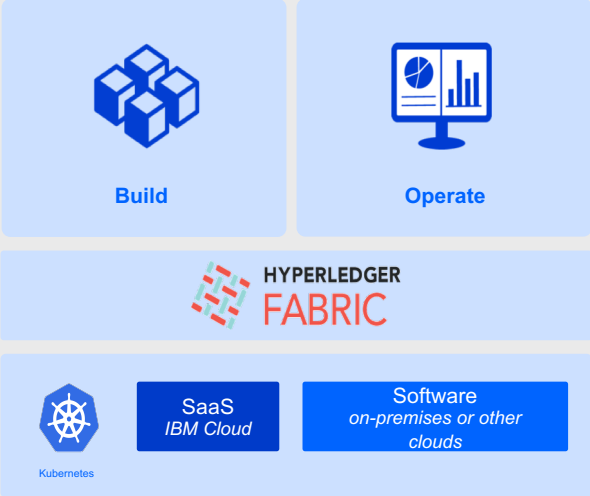
<http://ibm.biz/maciej-jedrzejczyk>

Quick Recap of IBM Blockchain

IBM Strategy



IBM Technology Stack



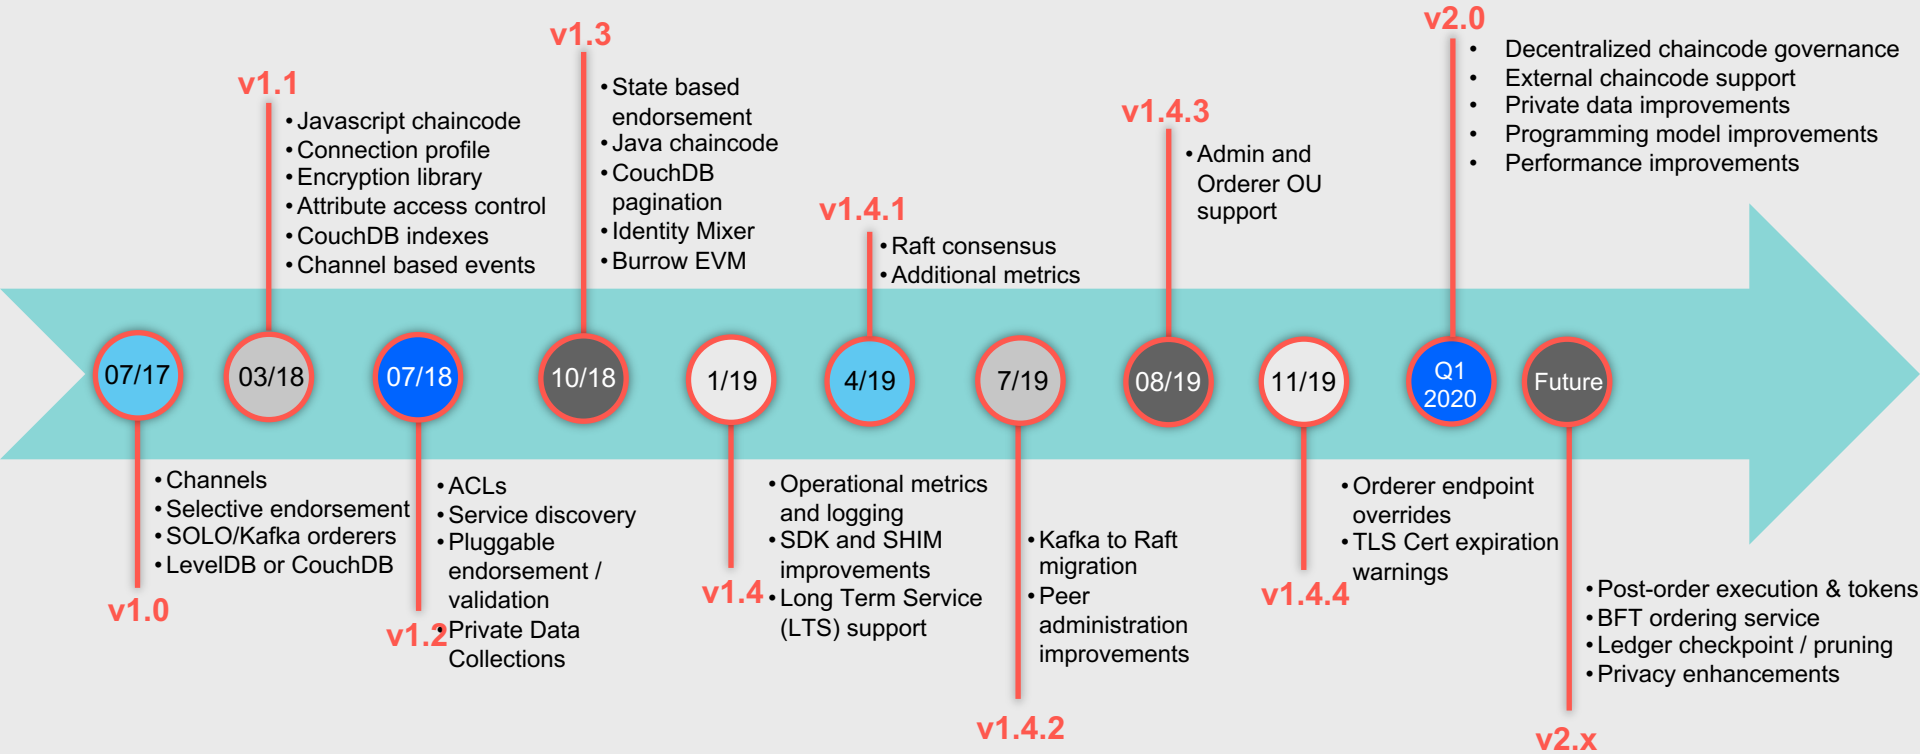
What is Hyperledger Fabric?



The image shows a screenshot of a web browser displaying the Hyperledger Fabric project page. The browser's address bar shows the URL <https://www.hyperledger.org/projects/fabric>. The page header includes the Linux Foundation logo and the Hyperledger logo. The main content area features the Hyperledger Fabric logo, a navigation menu with links for About, Members, Projects, Community, Industries, Resources, News & Events, and Blog, and social media icons. Below the navigation is a large blue banner with the text "HYPERLEDGER FABRIC" and two buttons: "GET THE CODE" and "BUILD YOUR FIRST NETWORK". A "Feedback" button is visible on the right side of the banner. Below the banner, the text "Type: DLT, Smart Contract Engine" and "Status: Active" is displayed. A video player titled "Hyperledger Fabric Explainer" is also visible. At the bottom of the page, the Hyperledger Fabric logo is repeated in a larger, stylized format.

- An implementation of blockchain technology that is a foundation for developing blockchain applications for business
- Emphasis on ledger, smart contracts, consensus, confidentiality, resiliency and scalability.
- V2.0.1 released February 2020
- V1.4 Long Term Service release with emphasis on production operational and serviceability enhancements; new programming model abstractions for ease of development
- IBM is one of the many contributing organizations

Hyperledger Fabric Roadmap



What's New in Hyperledger Fabric 2.0?

In this release, Hyperledger Fabric developers **focused on**:

- Chaincode lifecycle
- Private Data Collections enhancements
- Security and performance improvements

What this release **does not** introduce:

- Token management model
- BFT consensus
- Upgrades to Fabric-CA and zero-knowledge proofs

Warning! Hyperledger Fabric 2.0 **is not** an LTS release! LTS tag has not been dropped from 1.4.x yet.

What's new in chaincode

What's new in chaincode

Decentralized governance brings new patterns for collaboration and consensus

Multiple organizations must agree to the parameters of a chaincode

New channel policy (LifecycleEndorsement) which specifies the set of organizations required to approve a chaincode definition on a channel.

Committing a channel definition requires endorsement to fulfill LifecycleEndorsement policy

Number of organizations needed to approve definition is governed by the Channel/Application/LifecycleEndorsement policy

Chaincode definition includes: Endorsement policy, Private data collection configurations, etc

More: https://hyperledger-fabric.readthedocs.io/en/release-2.0/enable_cc_lifecycle.html#create-enable-lifecycle-json

```
"LifecycleEndorsement": {  
  "mod_policy": "Admins",  
  "policy": {  
    "type": 3,  
    "value": {  
      "rule": "MAJORITY",  
      "sub_policy": "Endorsement"  
    }  
  },  
  "version": "0"  
}
```

What's new in chaincode

New Chaincode Lifecycle Flow - Installation

Introduces a more deliberate chaincode installation and upgrade process

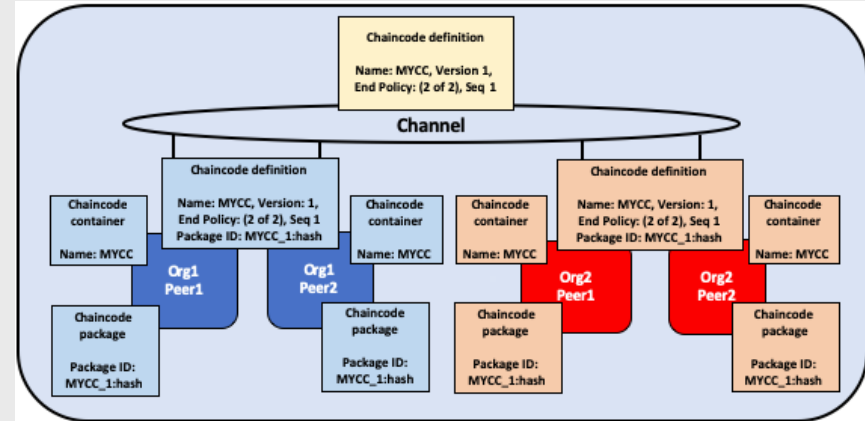
Each organization packages/installs chaincode to their peers

Individual organizations each call lifecycle system chaincode to approve a chaincode definition

```
$ peer lifecycle chaincode approveformyorg --init-required ...
```

When sufficient number of organizations have approved, one organization calls lifecycle system chaincode to commit the chaincode definition to the channel

```
$ peer lifecycle chaincode commit ...
```



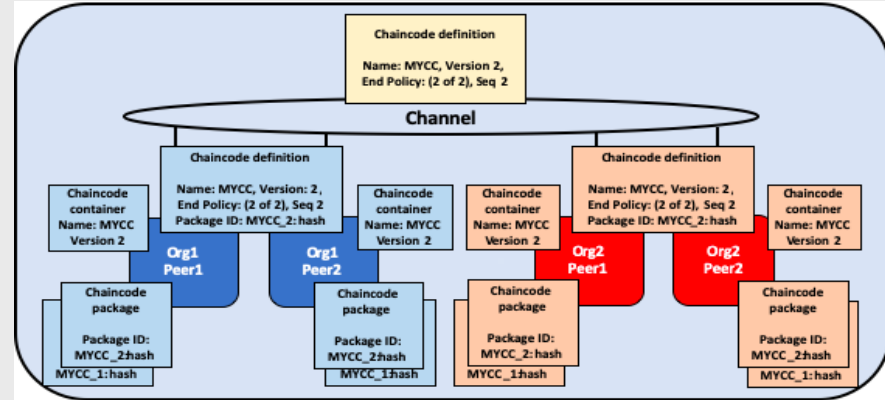
More: https://hyperledger-fabric.readthedocs.io/en/release-2.0/chaincode_lifecycle.html#upgrade-a-chaincode

What's new in chaincode

New Chaincode Lifecycle Flow - Upgrade

Chaincode can be upgraded using the same process as used to install for the first time

1. Repackage the chaincode (Only if chaincode updated)
2. Install the new chaincode package on peers (Only if chaincode updated)
3. Approve a new chaincode definition
4. Commit the definition to the channel
5. Upgrade the chaincode container (Automatic - Only if chaincode updated)



More: https://hyperledger-fabric.readthedocs.io/en/release-2.0/chaincode_lifecycle.html#upgrade-a-chaincode

New Chaincode Lifecycle Flow – Deployment Scenarios

Joining a channel

- New organization can start using chaincode after joining an existing channel
- No need to be commit definition again
- When using default LifecycleEndorsement policy, number of required approvals/endorsements will change automatically.

Updating an endorsement policy

- Can be done without repackaging/re-installing chaincode
- Will take effect after the new definition is committed
- No need to restart the chaincode container

Approving definition without installed chaincode

- Organizations can approve a chaincode definition without installing the chaincode package
- No need to provide a packageID as part of chaincode definition

More: https://hyperledger-fabric.readthedocs.io/en/release-2.0/chaincode_lifecycle.html#upgrade-a-chaincode

New Chaincode Lifecycle Flow – Conflict Scenarios

Organization disagrees on chaincode definition

- Organization cannot use chaincode unless definition is approved
- Organizations with different chaincode definition will not be able to execute the chaincode on their peers
- Other organizations in the channel can use chaincode for approved definition

Channel does not agree on a definition

- Happens when organizations approve definitions that do not match
- Definition cannot be committed to the channel
- None of the channel members will be able to use the chaincode

More: https://hyperledger-fabric.readthedocs.io/en/release-2.0/chaincode_lifecycle.html#upgrade-a-chaincode

What's new in chaincode

Chaincode packages are inspectable

Package in a tar file, ending with a .tar.gz file extension

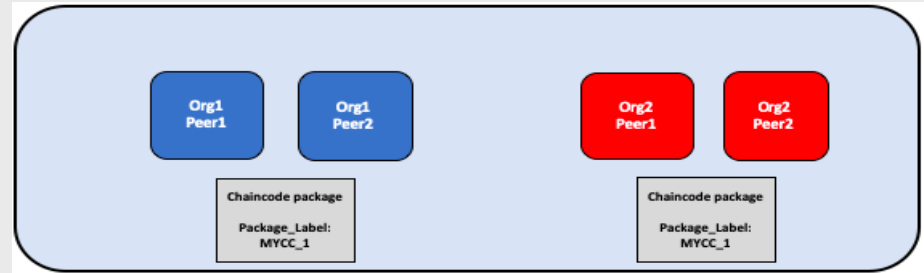
Two files: A metadata file “Chaincode-Package-Metadata.json” and tar containing the chaincode files

Metadata specifies the chaincode language, code path, and package label

It is not necessary for organizations to use the same package label

```
$ peer lifecycle chaincode package --label ...
```

More: <https://hyperledger-fabric.readthedocs.io/en/release-2.0/commands/peerchaincode.html#peer-chaincode-package>



Chaincode logic on peer and chaincode definition on channel are decoupled

The previous lifecycle defined each chaincode on the channel using a name and version that was specified when the chaincode package was installed. Chaincodes had to have identical logic and policies.

Multiple definitions can be used to run different instances of chaincode. Same chaincode can have multiple endorsement policies.

Ability to update chaincode definition (endorsement policy, collection configuration, etc) on the channel, without having to install new chaincode on each peer

Chaincode package and chaincode language installed on peer can now be different for each organization, so long as the ultimate chaincode execution results (RW sets) are the same on the required endorsers

Use Case #1: Each organization can now roll out minor fixes on their own schedules without requiring the entire network to proceed in lock-step.

Use Case #2: You can now use a single chaincode package and deploy it multiple times with different names on the same channel or on different channels to track different types of assets in their own chaincode namespace.

External chaincode support

Eliminate Docker daemon dependency – Previously peers were required to have access to a Docker daemon to build and launch chaincode, requiring the peer to have process privileges

Users can choose between running chaincode in a docker container (default) or using the external chaincode launcher feature to build and launch chaincode with the technology of their choice.

Chaincodes can now run as an external service (e.g. Kubernetes pod) which the peer connects to

Alternatives to containers – Chaincode no longer required to run in Docker containers, and instead can be run in the user's choice of environment (including containers)

External builder executables – A user provides a set of external builder executables to override how the peer prepares and launches chaincode

Location of external builder executables are configured in peer core.yaml.

More: https://hyperledger-fabric.readthedocs.io/en/release-2.0/cc_launcher.html

Other enhancements

Support for Node.js 12 and Java 11 for smart contracts
Moved to new Long Term Support release allowing smart contract developers to take advantage of the latest language features and performance enhancements
Investigating support for Node.js and Java SDKs

High-level programming model for Go smart contracts
Consistent experience across all supported languages, for smart contract developers who wish to use Go, including updated samples, tutorials, and documentation

More: <https://hyperledger-fabric.readthedocs.io/en/release-2.0/whatsnew.html>

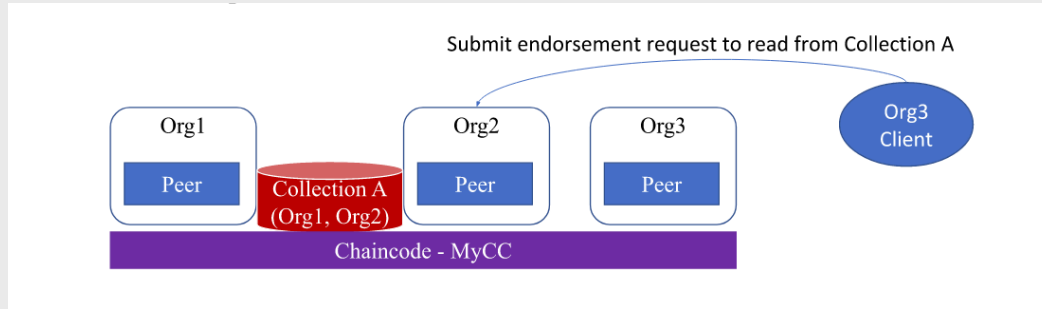
What's new in private data collections

What's new in private data collections

Read-Write access control on collections

Previously, any channel member could read-write to a collection if the chaincode did not have an adequate ACL

Now, the collection has `memberonlyread` and `memberonlywrite` attributes that can control that behaviour



```
{
  "name": "collectionMarblePrivateDetails",
  "policy": "OR('Org1MSP.member')",
  "requiredPeerCount": 0,
  "maxPeerCount": 3,
  "blockToLive": 3,
  "memberOnlyRead": true,
  "memberOnlyWrite": true,
  "endorsementPolicy": {
    "signaturePolicy": "OR('Org1MSP.member')"
  }
}
```

More: <http://www.bchainledger.com/2020/03/whats-new-in-hyperledger-fabric-v20.html>

What's new in private data collections

Sharing and verifying private data

Utilize implicit built-in private data collections per organization, no need to define

Each per-organization collection has an endorsement policy matching the organization
e.g. `PutPrivateData("_implicit_org_<MSPID>", key, value)`

Implicit per-organization private data collections allow for sharing private data at the key level on a need-to-know basis, e.g. when transferring an asset from one organization to another

Share and verify private data with new `GetPrivateDataHash()` chaincode API to verify on-chain hashes or when private data moves from one organization to another.



More: <http://www.bchainledger.com/2020/03/whats-new-in-hyperledger-fabric-v20.html>

What's new in private data collections

Collection level endorsement policy

Previously, any member could write to a collection if the transaction has adequate endorsement

Now it is possible to specify an endorsement policy for a private data collection

Overrides chaincode level endorsement policy for keys in the collection

Collections which organizations can write to a collection
Organizations can safely interact with smart contracts -
e.g. bid, approve, transfer workflows with data privacy and nonrepudiation

More: <https://hyperledger-fabric.readthedocs.io/en/release-2.0/endorsement-policies.html#setting-collection-level-endorsement-policies>

Security and performance enhancements

State database cache for improved performance

When using external CouchDB state database, read delays during endorsement and validation phases has historically been a performance bottleneck.

New peer cache replaces many of these expensive lookups with fast local cache reads.

The cache size can be configured by using the `core.yaml` property `cacheSize`.

Ordering Service Network performance improvements

Transaction validation is parallelized in the commit phase. The ordering node message processing flow has been optimized to remove redundant checks, and the write block processing is now asynchronous.

Ordering service nodes can now choose which channels it will serve instead of the legacy requirement of serving all channels.

Other Enhancements

Changes to the SDK, Fabric components and chaincode shim

Simplified wallet management for Node.js and Java SDKs provides a persistent data format that is more portable across different SDK language implementations, and pluggable persistent storage to allow users to make use of storage mechanisms appropriate to their deployment environment

The shim package and dependencies for go chaincode are no longer included in the chaincode build environment.

Chaincode that used the shim's `NewLogger()` will need to shift to a new logging mechanism.

The 'Solo' and 'Kafka' consensus type is deprecated.

Support for specifying orderer endpoints at the global level in channel configuration is deprecated.

Hyperledger Fabric Docker images will use Alpine Linux, a security-oriented, lightweight Linux distribution.

Rolling update enhancements

The upgrade documentation has been significantly expanded and reworked, and now have a standalone home in the documentation:

Upgrading to the latest release: <https://hyperledger-fabric.readthedocs.io/en/release-2.0/upgrade.html>

Upgrade considerations: https://hyperledger-fabric.readthedocs.io/en/release-2.0/upgrade_to_newest_version.html

Enabling the new chaincode lifecycle: https://hyperledger-fabric.readthedocs.io/en/release-2.0/enable_cc_lifecycle.html

Conclusions and next steps

Fabric Roadmap for 2.x in 2020

Major impact on the following personas:

- Stability & Resilience in Releases
- Tokens
- BFT Consensus
- Ledger Checkpointing
- Usability Updates

Follow JIRA for up to date information: <https://jira.hyperledger.org/projects/FAB/issues/FAB-16971?filter=allopenissues>

What do these changes mean for the business?

- An organization can add its own chaincode flavour and validations to better protect the organization's specific interests before chaincode is deployed or executed on a channel.
- Organizations can privately share data with other organizations in the channel on a need-to-know basis, eliminating the need to define channels or private data collections for many combinations of members

What do these changes mean for developers and operators?

- New topology models for Raft OSN to address performance and privacy issues.
- Changing approach to chaincode lifecycle and private data collections may increase flexibility and reduce operational overhead if automated checks are implemented.
- Implicit per-organization private data collections combined with collection level endorsement policies allow for a confidential transfer of assets without a need to implement ZKPs oraz encryption.
- Alpine-based images mean the only viable way to execute transactions on Peers and Orderers are through fabric-tools image.
- Using external chaincode launcher is tempting, especially when we consider a potential deployment on Kubernetes clusters.
- Operators can build and launch chaincode with the technology of their choice, removing the requirement to give the peer access to a Docker daemon.

Conclusion and next steps

Further Reading

Evolution of Fabric and the significance of the new release:

<https://www.hyperledger.org/blog/2020/01/30/welcome-hyperledger-fabric-2-0:-enterprise-dlt-for-production>

As always, you can find additional details of release content in the What's New documentation:

<https://hyperledger-fabric.readthedocs.io/en/release-2.0/whatsnew.html>

Information about Fabric changes and upgrade considerations can be found in the release notes:

<https://github.com/hyperledger/fabric/releases/tag/v2.0.0>

Download Hyperledger Fabric v2.0 today and give the new test network in the fabric-samples repository a try:

<https://hyperledger-fabric.readthedocs.io/en/release-2.0/install.html>

Great article going in-depth regarding Hyperledger Fabric 2.0:

<http://www.bchainledger.com/2020/03/whats-new-in-hyperledger-fabric-v20.html>

Fabric development roadmap:


[https://jira.hyperledger.org/browse/FAB-](https://jira.hyperledger.org/browse/FAB-13754?qjql=project%20%3D%20FAB%20AND%20issuetype%20%3D%20Epic%20AND%20fixVersion%20%3D%20v2.1.0)

[13754?qjql=project%20%3D%20FAB%20AND%20issuetype%20%3D%20Epic%20AND%20fixVersion%20%3D%20v2.1.](https://jira.hyperledger.org/browse/FAB-13754?qjql=project%20%3D%20FAB%20AND%20issuetype%20%3D%20Epic%20AND%20fixVersion%20%3D%20v2.1.0)

[0](https://jira.hyperledger.org/browse/FAB-13754?qjql=project%20%3D%20FAB%20AND%20issuetype%20%3D%20Epic%20AND%20fixVersion%20%3D%20v2.1.0)



Next Webinar: Deploying IBM Blockchain Platform on IBM Cloud



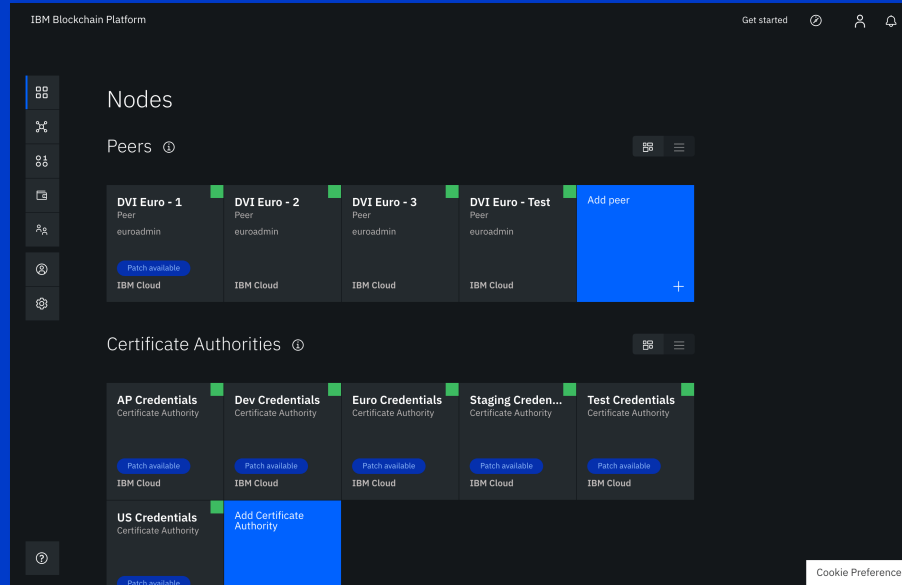
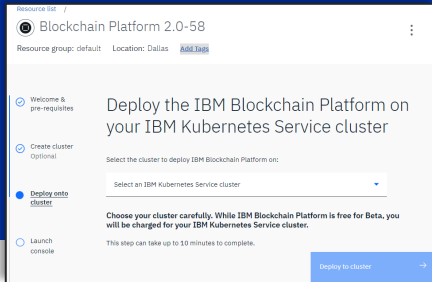
Blockchain Platform

IBM • IAM-enabled

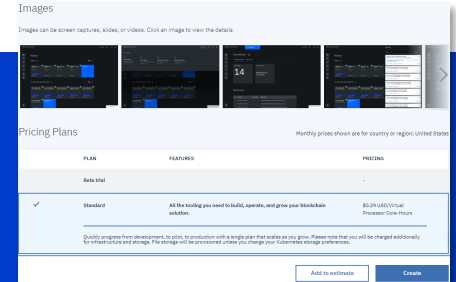
Welcome to the fast, flexible way to build, operate, and grow blockchain solutions.

1. Go to the IBM Cloud console (cloud.ibm.com) and search the catalog for blockchain

3. Create and attach an IBM Kubernetes Service cluster

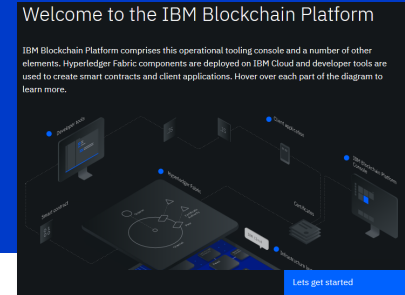


5. Success! Create or join networks, scale, customize, add HA...



2. Create your blockchain service instance

4. Launch blockchain console



IBM Blockchain

...and you're on your way!



Q&A