IBM Cloud

IBM Cloud Framework for Financial Services

Solution guide

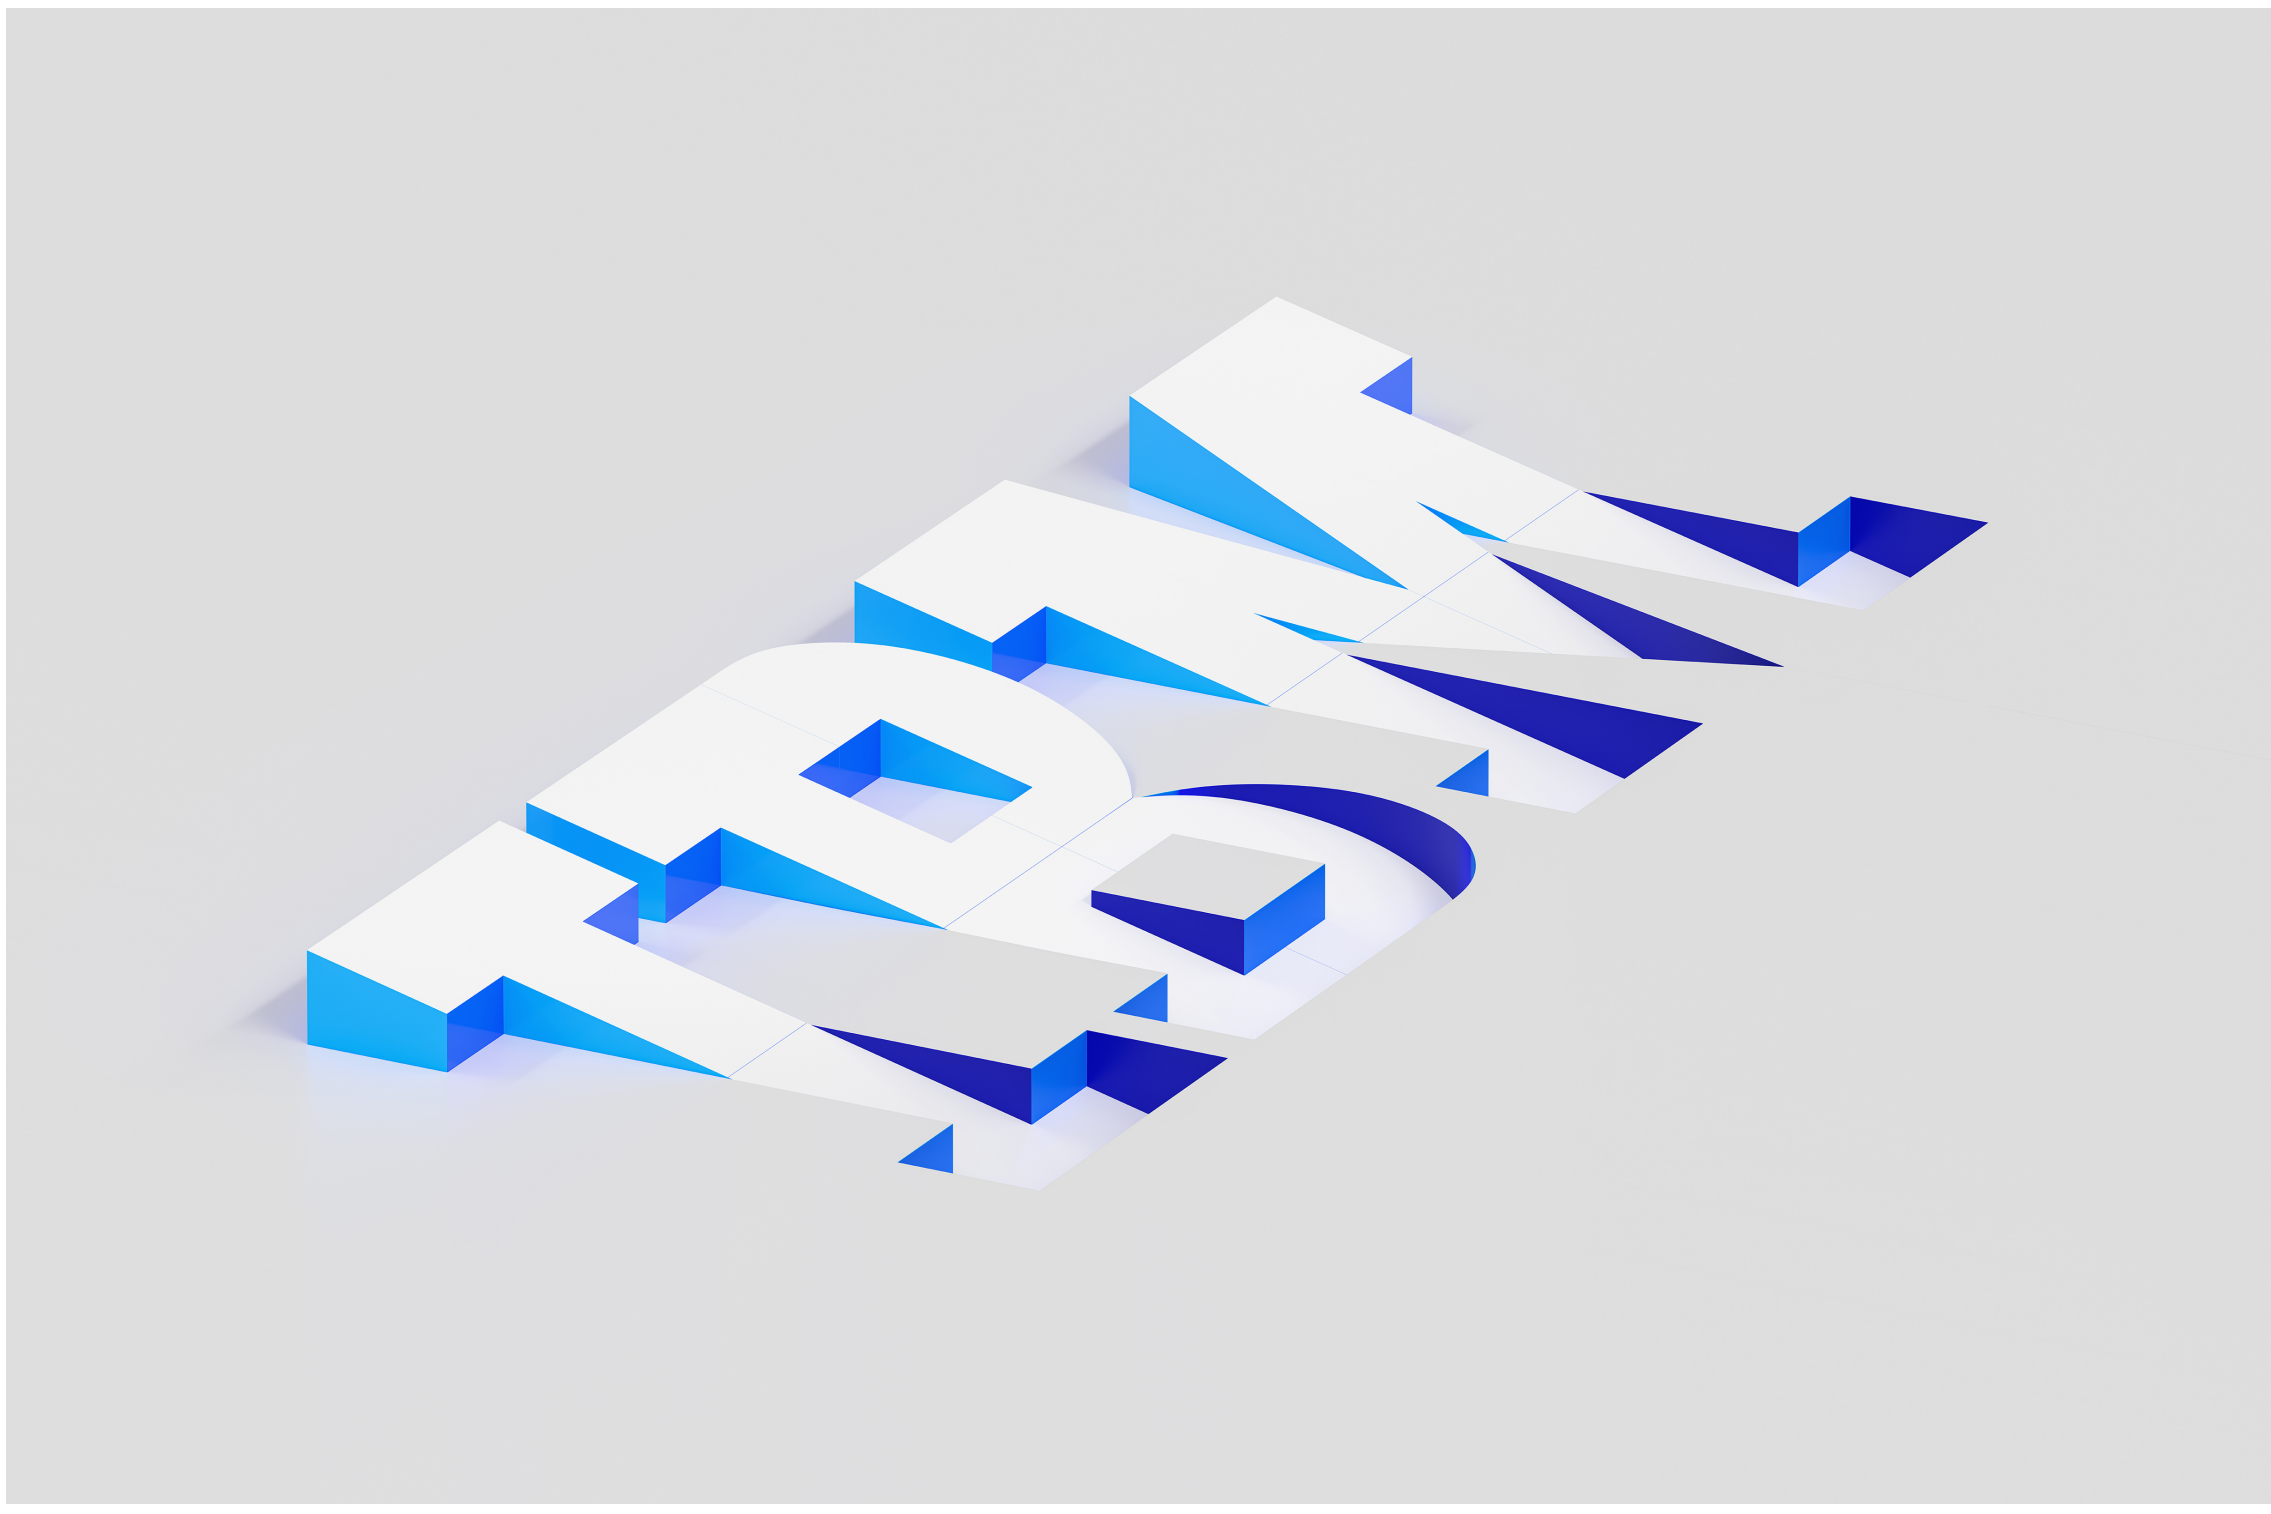

IBM

Edition notices

This PDF was created on 2026-07-30 as a supplement to *IBM Cloud Framework for Financial Services* in the IBM Cloud docs. It might not be a complete set of information or the latest version. For the latest information, see the IBM Cloud documentation at https://cloud.ibm.com/docs/framework-financial-services.

Getting started with IBM Cloud for Financial Services

IBM Cloud for Financial Services® is designed to build trust and enable a transparent public cloud ecosystem with the features for security, compliance, and resiliency that financial institutions require. Financial institutions can confidently host their mission-critical applications in the cloud and transact quickly and efficiently. With a large partner ecosystem of independent software vendors (ISVs), Software as a Service (SaaS), and fintech partners, IBM Cloud for Financial Services offers a new generation of cloud for the enterprise. Financial institutions can now deploy on public cloud to enable innovation and deliver new outstanding customer experiences, while managing stringent industry regulations for sensitive data and complex workloads.

IBM Cloud Framework for Financial Services

IBM Cloud Framework for Financial Services is designed to help address the needs of financial services institutions with regulatory compliance, security, and resiliency during the initial deployment phase and with ongoing operations. The framework also helps to simplify the ability of financial institutions to transact with ecosystem partners who deliver software or SaaS applications, and who meet the requirements of the framework.

The IBM Cloud Framework for Financial Services consists of:

- A comprehensive set of control requirements designed to help address the security requirements and regulatory compliance obligations of financial institutions and cloud best practices. The cloud best practices include a shared responsibility model across financial institutions, application providers, and IBM Cloud.
- Detailed control-by-control guidance for implementation and supporting evidence to help address the security and regulatory requirements of the financial industry.
- Reference architectures designed to facilitate compliance with the control requirements. In addition, resources are provided to deploy infrastructure as code in order to automate deployment and configuration of the reference architectures.
- Tools and IBM services, such as IBM Cloud® Compliance Manager, to enable parties to efficiently and effectively monitor compliance, remediate issues, and generate evidence of compliance.
- Ongoing governance of the framework documentation that considers new and changing regulations, as well as bank and public cloud requirements.

Control requirements

The technology-agnostic control requirements defined in the framework were built by the industry for the industry. The framework contains 608 control requirements that span 7 focus areas and 21 control families. The control requirements were initially based on NIST 800-53 and have been enhanced based on feedback from leading industry partners.

See Reference architectures have been defined which can be used as a basis for meeting the security and regulatory requirements defined by the controls. The reference architectures each center on specific IBM Cloud technologies, including:

- IBM Cloud® Virtual Private Cloud reference architecture, with options for using one or both of:
  - IBM Cloud® Virtual Servers for Virtual Private Cloud
  - Red Hat® OpenShift® on IBM Cloud®
- IBM Cloud Satellite® reference architecture
- IBM Cloud® for VMware® Regulated Workloads reference architecture

In addition, prescriptive deployment and configuration guidance is provided for the components in each of these reference architectures. The guidance is intended as an easier to consume supplement to the CIO templates that is more appropriate for architects and developers trying to build and implement applications for IBM Cloud for Financial Services.

Becoming IBM Cloud for Financial Services Validated

IBM designates IBM Cloud services as IBM Cloud for Financial Services Validated when the Services have been determined by IBM to materially implement IBM Cloud Framework for Financial Services control requirements. Similarly, IBM designates ecosystem partners as IBM Cloud for Financial Services Validated when the partners conform to the IBM Cloud for Financial Services® reference architectures and guidance and have been determined by IBM to materially implement the IBM Cloud Framework for Financial Services control requirements. Clients should review associated documentation made available by IBM and ecosystem partners to conduct their due diligence.

Through the shared responsibility model of the IBM Cloud Framework for Financial Services and the surrounding standardized processes, financial institutions and ecosystem partners get benefits such as:

- Less time that is spent by ecosystem partners demonstrating compliance and more time delivering innovative services.
- Reduction in the time and effort by financial institutions to ensure the compliance of third-party vendors, and more time spent delivering new, innovative services to their customers.
- Streamlined procurement, contracting, and onboarding within the ecosystem that leads to reduced time to market for all parties.

In addition, ecosystem partners (even those who are not yet IBM Cloud for Financial Services Validated) are encouraged to onboard to the IBM Cloud

catalog.

Next steps

If you plan to provide SaaS:

- Review Best practices for software as a service , which represent a generic set of implementation principles that apply to SaaS applications on both reference architectures.
- Explore the reference architectures and deployment guidance

If you plan to provide software that another organization will operate, review Best practices for software .

Security categorization, functionality and assurance

The reference architectures and associated guidance are intended to help you meet the Financial Services Validation process performs assurance activities to enable confidence in the implementation of the controls.

The overall security categorization of the data will guide the functionality and assurance required from a cloud service being used to protect the data. For example, while Hyper Protect Crypto Services and Key Protect have received the same level of assurance, both being Financial Services Validated, they each provide different security functionality. They may be used to protect data with a different security categorization. For a further discussion on the differences, read Data encryption at rest in IBM Cloud for a comparison of functionality.

When it comes to differences in assurance, there are multiple options for audit logging. You can use Activity Tracker Event Routing, that's Financial Services Validated, to forward logs to either IBM Cloud Object Storage or Event Streams (Enterprise Plan) and then use a custom solution for an audit event monitoring application. An alternative is to use Activity Tracker hosted event search, that's not Financial Services Validated, but has audit event monitoring built-in. The overall security categorization of the data should guide whether you need the assurance provided with a Financial Services Validated solution.

When you think about modifying a security solution because of a lower security categorization, think about the two dimensions of functionality and assurance.

Next steps

If you plan to provide SaaS:

- Review Best practices for software as a service , which represent a generic set of implementation principles that apply to SaaS applications on both reference architectures.
- Explore the reference architectures and deployment guidance

If you plan to provide software that another organization will operate, review Best practices for software .

Best practices and requirements

Best practices and requirements for software as a service

These best practices summarize some of the most important technical principles for SaaS providers based on the Control Implementation Overview templates, which are still the definitive guides to the controls and required evidence for application providers.

Tip: If you're a software provider, see Best practices for software.

1. Use an approved reference architecture

Requirement: Deploy and manage IBM Cloud resources as defined by an approved reference architecture.

Purpose & value: The quickest path toward meeting the controls of the IBM Cloud Framework for Financial Services is to use one of the pre-defined, single-tenant reference architectures. These architectures are designed to facilitate implementing the other requirements and best practices such as network isolation, security function isolation, high availability, backup and recovery, encryption, audit logging, compliance monitoring, etc.

A key aspect of the framework is to separate user workloads from system management functionality and isolate security functions from non-security functions. One of the most important features of the reference architectures is the definition of physical and logical separation between the edge, management, and workload planes. It is your responsibility to ensure that separation stays in place.

Implementation guidance: Explore and choose one of the supported reference architectures for the IBM Cloud for Financial Services:

- IBM Cloud® Virtual Private Cloud reference architecture, with options for using one or both of:
  - IBM Cloud® Virtual Servers for Virtual Private Cloud
  - Red Hat® OpenShift® on IBM Cloud®
- IBM Cloud Satellite® reference architecture
- IBM Cloud® for VMware® Regulated Workloads reference architecture

Most Relevant Controls:

Family	Control
System and Communications Protection (SC)	SC-2 Separation of System and User Functionality SC-3 Security Function Isolation

Related controls for use of reference architecture [FSv2.0]

Family	Control
System and Communications Protection (SC)	SC-2 Application Partitioning SC-3 Security Function Isolation

Related controls for use of reference architecture [FSv1.1]

2. Use only services that are IBM Cloud for Financial Services Validated

Requirement: Only use IBM Cloud and third-party managed services which are IBM Cloud for Financial Services Validated.

Purpose & value: IBM Cloud for Financial Services Validated designates that an IBM Cloud service or ecosystem partner service has evidenced compliance to the controls of the IBM Cloud Framework for Financial Services. Using services which are Financial Services Validated makes it much easier to meet all of the control requirements and provide evidence of compliance.

An external service is an IBM Cloud or third-party service deployed outside of the authorization boundaries of your offering and typically not under direct control of your offering. External services used by an offering should be considered extensions of the offering itself. For an offering to be validated against the requirements of the IBM Cloud Framework for Financial Services, all of its external services must meet the same requirements.

For cases where you need additional functionality not covered by any Financial Services Validated service, you are advised to install your own software within your deployment. However, you are then responsible for ensuring the software meets all of the IBM Cloud controls.

Important: If you represent a technology vendor seeking the Financial Services Validated designation, then you must request special approval for a temporary deviation if you wish to use services that are are not themselves Financial Services Validated. If you represent a financial institution or if the Financial Services Validated designation is not required by your consumers, then you have freedom to choose based on the level of risk you are willing to accept. Regardless of designation, all IBM Cloud services are designed with security in mind, and many are certified with other

compliance programs, such as ISO, SOC 2, etc.

Implementation guidance:

For technology vendors:

1. For services which are not Financial Services Validated, consult with your IBM customer success representative to determine if it is appropriate for use.
2. If you add services to your solution which are not Financial Services Validated, be sure to update the appropriate parts of the System Environment and Inventory section of your Control Implementation Overview template.

Most Relevant Controls:

Family	Control
Access Control (AC)	AC-20 Use of External Systems
System and Services Acquisition (SA)	SA-4 Acquisition Process SA-9 External System Services
Enterprise System and Services Acquisition (ESA)	ESA-5 Subcontractor Risk Management
Security Assessment and Authorization (CA)	CA-3 Information Exchange

Related controls for using only Financial Services Validated services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
System and Services Acquisition (SA)	SA-4 Acquisitions Process SA-9 External Information System Services
Enterprise System and Services Acquisition (ESA)	ESA-5 Subcontractor Risk Management
Security Assessment and Authorization (CA)	CA-3 System Interconnections

Related controls for using only Financial Services Validated services [FSv1.1]

3. Ensure all deployed software meets all controls

Requirement: Ensure all software used by your service is developed, installed, configured, and managed in accordance with the IBM Cloud Framework for Financial Services control requirements. It is recommended that you use software that is Financial Services Validated, when available.

Purpose & value: Your service will need software to function. This may be software you implement, or it may be software from third-parties, such as databases, logging stacks, bastion hosts, web application firewalls, etc. These are all perfectly acceptable if you adhere to all IBM Cloud Framework for Financial Services requirements related to development, deployment, configuration, and management of the software.

You will be expected to provide appropriate evidence for all software as part of the validation process. Using software that is already Financial Services Validated will make it easier for you to provide this evidence.

4. Implement a system of account, identity, and access management to enable a zero-trust environment

Requirement: Implement a system of zero-trust where account, identity, and access management are based on the principles of separation of duties and least privilege. This system must include IBM Cloud accounts and resources under the control of IBM Cloud Identity and Access Management (IAM) as well as supporting accounts and resources outside of IBM Cloud.

Purpose & value: The design of the IBM Cloud for Financial Services is intended to enable a zero-trust model for deployments of the reference architectures. The scope of any one individual operator or administrator is strictly limited to those actions necessary and appropriate to perform their assigned roles. Overlap of duties is minimized or eliminated to prevent the need for privilege escalation that may result in undesired flows of information into or out of your deployment. It is guided by two major principles:

- Separation of duties - No user should be given enough privileges to misuse the system on their own. For example, no one user should have the authority/ability to develop, compile, and/or move object code from non-production environments into production environments.

- Least privilege - Restrict the access privileges of authorized personnel (e.g., program execution privileges, file modification privileges) to the minimum necessary to perform their jobs. Access must be granted based only on a need to know. Access that has not been explicitly permitted is denied by default. Role-based access controls (RBAC) are used to allocate logical access to a specific job function or area of responsibility, rather than to an individual. Administrative access is only granted to individuals whose job roles have been approved for administration responsibilities.

Not only do these principles apply to IBM Cloud accounts and resources, but they also apply to accounts and resources that are not managed by IBM Cloud. For example, your operators will need credentials and proper authorizations to resources such as (but not limited to):

- Self-installed software (such as databases, firewalls, etc.)
- Source code control systems
- External identity providers and user directories
- etc.

In addition, you must provide RBAC for consumer access to the deployed application in the workload plane.

Individuals should be able to request access to resources, and they should be granted access only if it is absolutely required for the them to do their jobs. As individuals leave the organization, change job roles, etc., there needs to be a system to revalidate that the previously granted access continues to be needed for the individuals to perform their duties.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Access Control (AC)	AC-2 Account Management AC-3 Access Enforcement AC-5 Separation of Duties AC-6 Least Privilege AC-14 Permitted Actions Without Identification or Authentication
Identification and Authentication (IA)	IA-5 Authenticator Management IA-5 (1) Authenticator Management | Password-based Authentication

Related controls for zero trust [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management AC-3 Access Enforcement AC-5 Separation of Duties AC-6 Least Privilege AC-14 Permitted Actions without Identification or Authentication
Identification and Authentication (IA)	IA-5 Authenticator Management IA-5 (1) Authenticator Management | Password-Based Authentication

Related controls for zero trust [FSv1.1]

5. Use and maintain non-production environments for development and testing

Requirement: Do not put anything into production that has not been tested in an equivalent non-production environment. Put non-production environments in a separate IBM Cloud account. Treat all non-production environments as if they were for production environments by following all of the best practices and IBM Cloud Framework for Financial Services controls for each environment you create.

Purpose & value: Non-production environments (all environments that your service uses for development, testing, or demonstration) are as critical to the security of your service as production environments. Security weaknesses must be detected and mitigated in development environments by enforcing all the same controls and security testing as production. As security weaknesses in development environments may be more likely than production, separation of these environments from production is required.

Note: It is permissible to run a "pre-production" environment in your production account to be used only for final verification and sniff testing before promotion to production -- but, only if the set of operators is the same for both environments. However, all changes still need to be developed and tested first in a non-production environment in a separate account.

Implementation guidance:

1. Ensure all controls in the IBM Cloud Framework for Financial Services are followed for each new environment you create, whether for production or not.
2. Create and manage production environments using an IBM Cloud account separate from the one you are using for non-production environments.

Most Relevant Controls:

Family	Control
Configuration Management (CM)	CM-3 (2) Configuration Change Control | Testing, Validation, and Documentation of Changes CM-4 (1) Impact Analyses | Separate Test Environments
System and Services Acquisition (SA)	SA-3 (2) System Development Life Cycle | Use of Live or Operational Data SA-10 Developer Configuration Management
System and Communications Protection (SC)	SC-2 Separation of System and User Functionality SC-3 Security Function Isolation

Related controls for non-production environments [FSv2.0]

Family	Control
Configuration Management (CM)	CM-3 (2) Configuration Change Control | Testing, Validation, and Documentation Of Changes CM-4 (1) Impact Analyses | Separate Test Environments
System and Services Acquisition (SA)	SA-10 Developer Configuration Management SA-15 (9) Development Process, Standards, and Tools | Use of Live Data
System and Communications Protection (SC)	SC-2 Application Partitioning SC-3 Security Function Isolation

Related controls for non-production environments [FSv1.1]

6. Enforce information flow policies and protect the boundaries of your application

Requirement: Leverage the components of the reference architecture to enforce policies for information flows and to protect the boundaries of your application.

Purpose & value: Information flow control determines where information is allowed to travel within the components of your service offering as well as to/from other interconnected systems. Information flow control is intended to reduce the risk of sensitive information flowing from one system to another that has less stringent security policies. Flow control is based on the characteristics of the information (e.g., is the content of an information payload valid?) and/or the information path (e.g., is the information coming from a trusted IP address?).

Boundary protection increases security by monitoring and restricting communications at the external network boundary and key internal boundaries of your service. Boundary protection devices (e.g., gateways, encrypted tunnels, hardware and virtual firewalls, etc.) are commonly employed.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Access Control (AC)	AC-4 Information Flow Enforcement AC-4 (5) Information Flow Enforcement | Embedded Data Types AC-4 (6) Information Flow Enforcement | Metadata AC-4 (14) Information Flow Enforcement | Security or Privacy Policy Filter Constraints AC-4 (21) Information Flow Enforcement | Physical or Logical Separation of Information Flows AC-20 Use of External Systems
Security Assessment and Authorization (CA)	CA-3 Information Exchange

System and Communications Protection (SC)	SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls for boundary protection [FSv2.0]

Family	Control
Access Control (AC)	AC-4 Information Flow Enforcement AC-4 (5) Information Flow Enforcement | Embedded Data Types AC-4 (6) Information Flow Enforcement | Metadata AC-4 (14) Information Flow Enforcement | Security Policy Filter Constraints AC-4 (21) Information Flow Enforcement | Physical / Logical Separation of Information Flows AC-20 Use of External Information Systems
Security Assessment and Authorization (CA)	CA-3 System Interconnections
System and Communications Protection (SC)	SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls for boundary protection [FSv1.1]

7. Ensure all interactive operator actions are executed through a bastion host

Requirement: Ensure all interactive operator actions can only be run through a bastion host in your dedicated edge or management plane which is properly isolated from your workload plane. Enable recording of bastion sessions for auditing.

Purpose & value: A bastion host is a server in your edge or management plane that allows a secure connection to other resources in your edge, management, and workload planes. Using a bastion host for executing operator actions minimizes the chances of penetration from an external source. Session logging allows tracking of all activities that are performed through the bastion host in the event a potential security issue must be investigated.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Access Control (AC)	AC-6 (9) Least Privilege | Log Use of Privileged Functions AC-17 Remote Access
Audit and Accountability (AU)	AU-14 Session Audit
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (organizational Users) | Multi-factor Authentication to Privileged Accounts

Related controls for bastion host [FSv2.0]

Family	Control
Access Control (AC)	AC-6 (9) Least Privilege | Auditing Use of Privileged Functions AC-17 Remote Access

Audit and Accountability (AU)	AU-14 Session Audit
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (Organizational Users) | Network Access to Privileged Accounts

Related controls for bastion host [FSv1.1]

8. Capture audit events and forward to a SIEM

Requirement: Capture audit events for actions performed on the components of your service, including IBM Cloud services as well as any software components you install within your deployment. Audit logs should be securely stored, and they should also be forwarded to a security information and event management (SIEM) system.

Purpose & value: It is important to have audit logs to be able to investigate abnormal activity as well to comply with regulatory audit requirements. With a proper audit logging system, you can be automatically alerted to potential problems. This allows you to be proactive and check into potential security issues shortly after they've occurred.

Your software components (whether written by you or a third-party) should be enabled to emit critical management and data events. Many examples are listed in the AU-2 control from the IBM Cloud Framework for Financial Services, including:

- Access, downloading, revisions to confidential information
- Access to or update to any financial transaction data
- Any changes to consumer accounts
- Login/logoff event successes and failures
- etc.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Audit and Accountability (AU)	AU-1 Policy and Procedures AU-2 Event Logging AU-3 Content of Audit Records AU-11 Audit Record Retention
System and Information Integrity (SI)	SI-4 System Monitoring

Related controls for audit logs [FSv2.0]

Family	Control
Audit and Accountability (AU)	AU-1 Audit and Accountability Policy and Procedures) | Network Access to Privileged Accounts AU-2 Audit Events AU-3 Content of Audit Records AU-11 Audit Record Retention
System and Information Integrity (SI)	SI-4 Information System Monitoring

Related controls for audit logs [FSv1.1]

9. Ensure operational logging and monitoring is implemented

Requirement: Ensure operational logging and monitoring is implemented.

Purpose & value: Operational logging and operational monitoring are a very important for ensuring your application runs smoothly. Proper operational logging and monitoring can help you determine if you need to failover to an alternate storage or processing site. In addition, operational logging and monitoring can help you determine if operations have returned to normal after a system disruption.

Operational logging is a complement to audit logs. Audit logs contain auditable events, while operational logs contain everything else that might be logged. For example, an operational log might contain entries that reflect what's happening as a program executes, such as functions getting called, stack traces getting thrown, etc. While this kind of log data is key to keeping a system running smoothly, it's typically not sent to a SIEM.

Operational logs can be split into two categories:

- Application - Log data generated by software components that you deploy and manage within your deployment. This may be from your own code, or other software components like databases, message queues, etc.
- Platform - Log data from IBM Cloud service instances that is not contained in the audit data sent to Activity Tracker Event Routing.

Operational monitoring for gauging system health is a very important complement to monitoring for security and compliance. Operational metrics include measurements for CPU usage, memory usage, API response times, etc.

Most Relevant Controls:

Family	Control
Contingency Planning (CP)	CP-2 (3) Contingency Plan | Resume Mission and Business Functions CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-10 System Recovery and Reconstitution
System and Information Integrity (SI)	SI-11 Error Handling

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Contingency Planning (CP)	CP-2 (3) Contingency Plan | Resume Essential Missions / Business Functions CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-10 Information System Recovery and Reconstitution
System and Information Integrity (SI)	SI-11 Error Handling

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

10. Follow secure development processes and ensure software integrity

Requirement: Ensure that security engineering principles are applied in the design, development, implementation, and modification of the system. Maintain software integrity by using signed images, applying security patches, doing vulnerability scans, etc.

Purpose & value: It is critical to follow security engineering principles and manage changes to your service using a system development life cycle (SDLC) that incorporates information security considerations. Doing so minimizes the risk of introducing code or configuration changes that undermine your system security.

Lack of software integrity leaves you vulnerable to security problems. So, it is critical for you to be able to assert the origin of all software components in the system and to keep software up to date. In addition, running regular scans on software to detect vulnerabilities allows you to more quickly take corrective action.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
System and Information Integrity (SI)	SI-7 Software, Firmware, and Information Integrity
System and Services Acquisition (SA)	SA-3 System Development Life Cycle SA-3 (2) System Development Life Cycle | Use of Live or Operational Data SA-8 Security and Privacy Engineering Principles SA-10 Developer Configuration Management SA-10 (1) Developer Configuration Management | Software and Firmware Integrity Verification SA-11 Developer Testing and Evaluation SA-15 Development Process, Standards, and Tools
Risk Assessment (RA)	RA-5 Vulnerability Monitoring and Scanning

Related controls for development processes [FSv2.0]

Family	Control
System and Information Integrity (SI)	SI-7 Software & Information Integrity
System and Services Acquisition (SA)	SA-3 System Development Life Cycle SA-8 Security Engineering Principles SA-10 Developer Configuration Management SA-10 (1) Developer Configuration Management | Software and Firmware Integrity Verification SA-11 Developer Security Testing and Evaluation SA-15 Development Process, Standards, and Tools SA-15 (9) Development Process, Standards, and Tools | Use of Live Data
Risk Assessment (RA)	RA-5 Vulnerability Scanning

Related controls for development processes [FSv1.1]

11. Encrypt consumer data at rest and in transit

Requirement: Encrypt all consumer data whether at rest or in transit.

Purpose & value: Consumers do not want their data to be accessible to bad actors. Protecting consumer data against unauthorized disclosure, modification, or destruction throughout the data lifecycle is of paramount importance in the IBM Cloud for Financial Services. Cryptographic controls must be in place in all regions and availability zones to protect the confidentiality and integrity of data.

Data at rest is to always be encrypted, and you must use Hyper Protect Crypto Services to manage encryption keys. Hyper Protect Crypto Services allows you to take the ownership of the cloud HSM to fully manage your encryption keys and to perform cryptographic operations using Keep Your Own Key (KYOK). This means you have full control and authority over encryption keys. No one except you (not even IBM) has access to your encryption keys. You should ensure that all Financial Services Validates services which support integration with Hyper Protect Crypto Services are properly configured.

Data in transit should always be encrypted using TLS 1.2 or higher. This includes *all* traffic in your deployment such as between virtual server instances and between pods in Red Hat OpenShift on IBM Cloud.

Use Hyper Protect Crypto Services for TLS offload for all data that is that is requested from outside of IBM Cloud (inbound traffic to IBM Cloud) and protected by a certificate which is signed by a public Certificate Authority. Configure any web servers to use TLS offload to set up the session, so that the private key never leaves Hyper Protect Crypto Services.

More detailed guidance can be found in the Cryptographic Requirements appendix of the Control Implementation Overview Template for whichever reference architecture you are using.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
System and Communications Protection (SC)	SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection SC-28 Protection of Information at Rest SC-28 (1) Protection of Information at Rest | Cryptographic Protection

Related controls for data encryption [FSv2.0]

Family	Control

System and Communications Protection (SC)	SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection SC-28 Protection of Information At Rest SC-28 (1) Protection of Information at Rest | Cryptographic Protection

Related controls for data encryption [FSv1.1]

12. Implement business continuity and disaster recovery

Requirement: Implement a system for business continuity and disaster recovery (BCDR). Define Recovery Point Objective (RPO), Recovery Time Objective (RTO), and Maximum Tolerable Downtime (MTD) metrics for essential business functions. Implement your application so that the metrics are achieved.

Purpose & value: Unfortunately, disasters occur that can make some or all of the data centers in a region unavailable. To avoid data loss, you must ensure all important data is backed up to a separate multizone region that can be used to restore service. However, your consumers must be able to opt out of having their data stored in the alternate region based on their data residency requirements.

In addition, you must follow best practices for BCDR as defined by the specific managed services that you use. For example, Direct Link, Hyper Protect Crypto Services, etc.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Contingency Planning (CP)	CP-2 Contingency Plan CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-9 System Backup CP-10 System Recovery and Reconstitution CP-10 (2) Information System Recovery and Reconstitution | Transaction Recovery

Related controls for business continuity and disaster recovery [FSv2.0]

Family	Control
Contingency Planning (CP)	CP-2 Contingency Plan CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-9 Information System Backup CP-10 Information System Recovery and Reconstitution CP-10 (2) System Recovery and Reconstitution | Transaction Recovery

Related controls for business continuity and disaster recovery [FSv1.1]

13. Design your application for high availability (recommended)

Requirement: Design your application for high availability (HA) to minimize downtime for your consumers.

> **Note:** This item is unique among the other best practices in that it represents a recommendation rather than a IBM Cloud Framework for Financial Services requirement. The IBM Cloud Framework for Financial Services only requires you to meet the RPO, RTO, and MTD goals that you define (see previous best practice). However, to be best in class, you should strive to be highly available.

Purpose & value: Consumers will be relying on your service for critical operations and have the potential to be deeply impacted by any downtime. To avoid downtime, you should implement redundancy to eliminate single points of failure. It is strongly recommended that you deploy your application to:

- Multiple availability zones within any multizone region you're using. This enables you to ensure if one data center in a region goes down, that your application can continue to function.
- Multiple multizone regions with failover between them. If something catastrophic happens that impacts an entire region, you can failover to another

so consumers can still use your service.

In addition, it is recommended that you:

- Implement a solution for autoscaling your deployment before capacity is exceeded.
- Follow best practices for HA as defined by the specific managed services that you use. For example, Direct Link, Hyper Protect Crypto Services, etc.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Contingency Planning (CP)	CP-2 Contingency Plan CP-7 Alternate Processing Site
System and Communications Protection (SC)	SC-6 Resource Availability

Related controls for high-availability [FSv2.0]

Family	Control
Contingency Planning (CP)	CP-2 Contingency Plan CP-7 Alternate Processing Site
System and Communications Protection (SC)	SC-6 Resource Availability

Related controls for high-availability [FSv1.1]

14. Use endpoint detection and remediation (EDR) tooling to detect malicious code

Requirement: Use endpoint detection and remediation (EDR) tooling that is effective at detecting malicious code in a cloud environment.

Purpose & value: Malicious code includes viruses, worms, Trojan horses, and spyware. The presence of malicious code puts the operation of your service and the customer data you process at risk. It is important to run regular, automated scans of the components of your service to detect these vulnerabilities so that they can be mitigated.

There are many EDR solutions for virtual server instances such as CrowdStrike and Carbon Black. For Kubernetes, relevant tooling includes StackRox, NeuVector, Sysdig Secure, and Twistlock.

Most Relevant Controls:

Family	Control
System and Information Integrity (SI)	SI-3 Malicious Code Protection

Related controls for business continuity and disaster recovery [FSv2.0]

Family	Control
System and Information Integrity (SI)	SI-3 Malicious Code Protection

Related controls for business continuity and disaster recovery [FSv1.1]

15. Regularly scan for open ports / protocols

Requirement: Regularly scan for open ports / protocols, and ensure only those ports / protocols needed by your service are open.

Purpose & value: As part of the principal of least privilege, you should only open the minimum number of ports / protocols that are needed by your service. Open ports / protocols not needed by your service enable unnecessary threat vectors. So, it is important to regularly scan to make sure the list of open ports / protocols is correct. Tools such as Nmap or Tenable can be used for this purpose.

Most Relevant Controls:

Family	Control
Configuration Management (CM)	CM-7 Least Functionality CM-7 (1) Least Functionality | Periodic Review

Related controls for port scanning [FSv2.0]

Family	Control
Configuration Management (CM)	CM-7 Least Functionality CM-7 (1)Least Functionality | Periodic Review

Related controls for port scanning [FSv1.1]

16. Secure and manage secrets and certificates

Requirement: Secrets must be securely protected through their entire lifecycle.

Purpose & value: A secret is any piece of data that is sensitive within the context of an application or service. Secrets include all of the following but are not limited to:

- Passwords of any type (database logins, OS accounts, functional IDs, etc.)
- API keys
- Long-lived authentication tokens (OAuth2, github, IAM, etc.)
- SSH keys
- Encryption keys
- Other private keys (PKI/TLS certificates, HMAC keys, signing keys, etc.)

You should should ensure:

- Secrets are generated and stored in the environment (e.g., dev, test, production) where your service is deployed.
- Secrets never leave their environments (e.g., dev, test, production) and should be secured using access control measures. Service design should minimize the number of machines and people with access to secrets using both authorization and network restrictions based on the principle of least privilege.
- Secrets are rotated in according with the requirements of the IBM Cloud Framework for Financial Services with minimal or no down time.
- Secrets are never stored in source code, configuration files, or documentation.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Access Control (AC)	AC-2 Account Management
Identification and Authentication (IA)	IA-2 Identification and Authentication (organizational Users) IA-3 Device Identification and Authentication IA-5 Authenticator Management

Related controls for secrets management [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management
Identification and Authentication (IA)	IA-2 User Identification and Authentication IA-3 Device Identification and Authentication IA-5 Authenticator Management

Related controls for secrets management [FSv1.1]

17. Tag all IBM Cloud resources with security attributes

Requirement: Tag all IBM Cloud resources with security attributes.

Purpose & value: All IBM Cloud resources in your account should be tagged with security attributes you define. You can apply user tags to organize your resources and easily find them later, to help you with identifying specific team usage or cost allocation, and to control access to your resources without requiring updates to your IAM policies.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Access Control (AC)	AC-16 Security and Privacy Attributes
System and Communications Protection (SC)	SC-16 Transmission of Security and Privacy Attributes

Related controls for non-production environments [FSv2.0]

Family	Control
Access Control (AC)	AC-16 Security Attributes
System and Communications Protection (SC)	SC-16 Transmission of Security Attributes

Related controls for non-production environments [FSv1.1]

18. Monitor for security and compliance against a baseline configuration

Requirement: Deploy and use tooling for monitoring and reporting of security and compliance. Maintain a baseline configuration for your service that consists of automated mechanisms to facilitate information system baseline management.

With Compliance Manager you can embed security checks into your every day workflows to help monitor for security and compliance. By monitoring for risks, you can identify security vulnerabilities and quickly work to mitigate the impact and fix the issue. By using Compliance Manager along with external integrations (such as, OpenShift Compliance Operator (OSCO), Tanium, NeuVector, and so on), you can build a robust approach for monitoring for security and compliance issues.

Purpose & value: Continuous monitoring and reporting helps detect malicious activity and security vulnerabilities as early as possible. Automated baseline management makes it much more likely that any compliance-breaking changes to your system configurations are caught early.

Implementation guidance:

- VPC reference architecture
- Satellite reference architecture

Most Relevant Controls:

Family	Control
Change Management (CM)	CM-2 (2) Baseline Configuration | Automation Support for Accuracy and Currency CM-6 Configuration Settings CM-6 (1) Configuration Settings | Automated Management, Application, and Verification

Related controls for monitoring security and compliance [FSv2.0]

Family	Control
Change Management (CM)	CM-2 (2) Baseline Configuration | Automation Support for Accuracy and Currency CM-6 Configuration Settings CM-6 (1) Configuration Settings | Automated Management, Application, and Verification

Related controls for monitoring security and compliance [FSv1.1]

Next steps

Learn about the three reference architectures for the IBM Cloud for Financial Services:

- VPC reference architecture
- Satellite reference architecture
- VMware Regulated Workloads reference architecture

Or, if you're intending to be a software provider, see also:

- Best practices for software

Best practices and requirements for software

These best practices summarize some of the most important technical principles for software providers based on the controls and implementation guidance of the IBM Cloud Framework for Financial Services. These best practices are *not* a replacement for the information provided in the Control Implementation Overview templates, which are still the definitive guides to the controls and required evidence for application providers.

Tip: If you're a SaaS provider, see Best practices for software as a service.

1. Enable operators to properly configure your application to meet the controls

Requirement: Enable organization operating your software to meet the controls of the IBM Cloud Framework for Financial Services.

Purpose & value: An operator is responsible for the IBM Cloud Framework for Financial Services controls for every piece of software they run in their management or workload VPCs, whether running in a virtual server instance or in Red Hat OpenShift on IBM Cloud containers. So, the overarching principle for a software provider is to enable their software to be deployed, configured, and managed in a way that could be Financial Services Validated.

Implementation guidance: Provide clear written instructions for an operator to follow in order to deploy, configure, and manage your software on the IBM Cloud for Financial Services.

2. Provide account and access management to enable a zero-trust environment

Requirement: Enable operators to implement a system of zero-trust where account, identity, and access management is based on the principles of separation of duties and least privilege.

Purpose & value: The design of the IBM Cloud for Financial Services is intended to enable a zero-trust model for deployments of the reference architectures. The scope of any one individual operator or administrator is strictly limited to those actions necessary and appropriate to perform their assigned roles. Overlap of duties is minimized or eliminated to prevent the need for privilege escalation that may result in undesired flows of information into or out of the operator's deployment. It is guided by two major principles:

- Separation of duties - No user should be given enough privileges to misuse the system on their own. For example, no one user should have the authority/ability to develop, compile, and/or move object code from non-production environments into production environments.
- Least privilege - Restrict the access privileges of authorized personnel (e.g., program execution privileges, file modification privileges) to the minimum necessary to perform their jobs. Access must be granted based only on a need to know. Access that has not been explicitly permitted is denied by default. Role-based access controls (RBAC) are used to allocate logical access to a specific job function or area of responsibility, rather than to an individual. Administrative access is only granted to individuals whose job roles have been approved for administration responsibilities.

Implementation guidance: As a software provider, you must provide facilities to accomplish these goals as appropriate. As a simple example, if your software provides a database, you should provide the ability to create database users with different roles (e.g., viewer, editor, administrator).

Ideally, you should support integration with an external (to the application) identity and access management provider for validating users and their associated access group membership.

Most Relevant Control Groups or Controls:

Family	Control
Access Control (AC)	AC-2 Account Management AC-3 Access Enforcement AC-5 Separation of Duties AC-6 Least Privilege AC-14 Permitted Actions Without Identification or Authentication

Related controls for zero trust [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management AC-3 Access Enforcement AC-5 Separation of Duties AC-6 Least Privilege AC-14 Permitted Actions without Identification or Authentication

Related controls for zero trust [FSv1.1]

3. Enable ability to capture audit events

Requirement: Provide technical capability within the application to capture audit events.

Purpose & value: It is important that the operator has access to audit logs to be able to investigate abnormal activity as well to comply with regulatory audit requirements. With a proper audit logging system, an operator can be automatically alerted to potential problems. This allows the operator to be proactive and check into potential security issues shortly after they've occurred.

Implementation guidance: You must log critical events related to administering your software that can be accessed by the operator. Many examples of auditable events are listed in the AU-2 control from the IBM Cloud Framework for Financial Services, including:

- Access, downloading, revisions to confidential information
- Access to or update to any financial transaction data
- Any changes to consumer accounts
- Login/logoff event successes and failures
- etc.

The audit events that are generated should comply with the Cloud Auditing Data Federation (CADF) standard using the system time for timestamps.

The software should buffer audit logs locally in the event that the central audit facility is unavailable, and if the capacity for buffering the audit records is exceeded and audit events are discarded, the developer should send an audit event indicating the loss of audit records once the central audit facility is accessible.

You must provide the technical capability to ensure only permitted users have access to audit records.

Most Relevant Control Groups or Controls:

Family	Control
Audit and Accountability (AU)	AU-2 Event Logging AU-3 Content of Audit Records AU-5 Response to Audit Logging Process Failures AU-8 Time Stamps AU-9 Protection of Audit Information AU-10 Non-repudiation AU-11 Audit Record Retention AU-12 Audit Record Generation

Related controls for audit logs [FSv2.0]

Family	Control
Audit and Accountability (AU)	AU-2 Audit Events AU-3 Content of Audit Records AU-5 Response to Audit Processing Failures AU-8 Time Stamps AU-9 Protection of Audit Information AU-10 Non-repudiation AU-11 Audit Record Retention AU-12 Audit Generation

Related controls for audit logs [FSv1.1]

4. Follow secure development processes and ensure software integrity

Requirement: Ensure that security engineering principles are applied in the design, development, implementation, and modification of your software. Maintain software integrity by using signed images, applying security patches, doing vulnerability scans, etc.

Purpose & value: It is critical to follow security engineering principles and manage changes to your software using a system development life cycle (SDLC) that incorporates information security considerations. Doing so minimizes the risk of introducing code or configuration changes that undermine your software's security.

Lack of software integrity leaves your operator vulnerable to security problems. So, it is critical for you to be able to assert the origin of all software components in your application and to keep software components up to date.

Most Relevant Control Groups or Controls:

Family	Control
Risk Assessment (RA)	RA-9 Criticality Analysis
System and Services Acquisition (SA)	SA-8 Security and Privacy Engineering Principles SA-9 External System Services SA-11 Developer Testing and Evaluation SA-15 Development Process, Standards, and Tools

Related controls for development processes [FSv2.0]

Family	Control
System and Services Acquisition (SA)	SA-8 Security Engineering Principles SA-11 Developer Security Testing and Evaluation SA-12 Supply Chain Protection SA-14 Criticality Analysis SA-15 Development Process, Standards, and Tools

Related controls for development processes [FSv1.1]

5. Ensure consumer data can be encrypted at rest and in transit

Requirement: Ensure operator can configure your software so that consumer data can be encrypted properly at rest and in transit.

Purpose & value: Consumers do not want their data to be accessible to bad actors. Protecting consumer data against unauthorized disclosure, modification, or destruction throughout the data lifecycle is of paramount importance in the IBM Cloud for Financial Services. Encryption is a major component of this protection, and operators must be able to configure your software to meet the data encryption requirements of the IBM Cloud Framework for Financial Services.

Implementation guidance:

- You must implement and document application capabilities and configuration requirements in accordance with the cryptographic specifications provided in the Supplemental Guidance for Cryptography section of the CIO template.
- You must implement and document application capabilities and configuration requirements related to using cryptographic keys provided via Hyper Protect Crypto Services to protect data in transit and at rest.

Most Relevant Control Groups or Controls:

Family	Control
System and Communications Protection (SC)	SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection SC-28 Protection of Information at Rest SC-28 (1) Protection of Information at Rest | Cryptographic Protection

Related controls for data encryption [FSv2.0]

Family	Control

System and Communications Protection (SC)	SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection SC-28 Protection of Information At Rest SC-28 (1) Protection of Information at Rest | Cryptographic Protection

Related controls for data encryption [FSv1.1]

6. Enable software to be run in a highly available manner

Requirement: Design your software so that operators can meet the IBM Cloud Framework for Financial Services requirements for high availability.

Purpose & value: Operators SaaS offerings are required to be highly available so that consumers do not face downtime of critical applications. Typically, high availability is addressed mostly by the operator, but you will need to provide documentation if your software or software deployment pattern includes enhanced availability capabilities (such as replication, active-active, active-passive, etc.).

Most Relevant Control Groups or Controls:

Family	Control
System and Communications Protection (SC)	SC-6 Resource Availability

Related controls for high-availability [FSv2.0]

Family	Control
System and Communications Protection (SC)	SC-6 Resource Availability

Related controls for high-availability [FSv1.1]

Next steps

If you haven't already, you should learn about the three reference architectures for the IBM Cloud for Financial Services where your software would be deployed:

- IBM Cloud® Virtual Private Cloud reference architecture, with options for using one or both of:
  - IBM Cloud® Virtual Servers for Virtual Private Cloud
  - Red Hat® OpenShift® on IBM Cloud®
- IBM Cloud Satellite® reference architecture
- IBM Cloud® for VMware® Regulated Workloads reference architecture

Release notes

Use the release notes to learn about the latest changes to the documentation for the IBM Cloud for Financial Services®.

Note: Looking for IBM Cloud status, platform release notes, security bulletins, or maintenance notifications? See IBM Cloud status.

23 June 2026

- Added guidance for Cloud Logs best practices

  Learn to use IBM® Cloud Logs features while maintaining Financial Services compliance and reducing costs.

17 June 2026

- Added guidance for data loss prevention and perimeter acccess patterns

  Learn about the data loss prevention strategies with a focus on perimeter access patterns for Independent Software Vendor (ISV) deployments in IBM Cloud for Financial Services, including security enforcements, architectural decisions, and threat mitigations.

15 June 2026

- Added guidance for Direct Link for on prem connectivity

  Direct Link can provide a secure and dedicated netowork connection between on-prem environment ("Bank") and ISV Cloud account for the ISV to access on-prem services (Bank information systems).

- Added link to multi-tenant solution guidance

  The guide provides different deployment models focused mostly on multitenancy for container-based solutions.

10 June 2026

- Added guidance for egress proxy

  FS controls require SG and ACL to not have 0.0.0.0/0 for egress rules. In high security environments, having an egress for 0.0.0.0/0 is not acceptable. The solution architecture must provide a way to access public services on the internet which do not publish IP Address from isolated secured environments on cloud and still be able to meet FS controls with minimal exceptions.

2 June 2026

- Added guidance for container image management

  This pattern presents a set of architectural alternatives for controlling access to container images used in workload deployments on Red Hat OpenShift on IBM Cloud clusters. The result is an approach that enables secure workload deployment with required controls over software supply chain elements.

24 March 2025

- Updated IBM Cloud Framework with changes for FS 2.0.

30 April 2023

- Updated Infrastructure as Code (IaC) article to point at new deployable architectures.

24 April 2023

- Removed recommendation that limited deployments to a subset of available multizone regions.

31 March 2023

- Integrated references to supplemental guidance for enterprise account architecture for large enterprises.

03 March 2023

- Updated list of multizone regions which are IBM Cloud for Financial Services Validated to include Toronto (`ca-tor`) and Sydney (`au-syd`).

23 January 2023

- Updated list of services which are designated as IBM Cloud for Financial Services Validated to include IBM Cloud Compliance Manager, IBM Cloud® Continuous Delivery, IBM Cloud® Secrets Manager, and IBM Cloud® Client VPN for VPC.

30 September 2022

- Added contextual links into IBM Cloud Satellite® reference architecture as well as deployment and configuration guidance for that architecture. During this process, there were numerous organizational changes that were made due to the overlap between the IBM Cloud Satellite® reference architecture and the IBM Cloud VPC reference architecture.
- Enabled link to spreadsheet of control requirements in IBM Cloud Framework for Financial Services.
- Added article Deploy infrastructure as code for reference architectures about Terraform automation for IBM Cloud for Financial Services.
- Enhanced content on IBM Cloud® Compliance Manager and IBM Cloud for Financial Services profile. See Compliance monitoring.
- Added site map.
- Made many editorial updates.

17 March 2022

- Clarify guidance around using IBM Cloud Financial Services Validated services depending on whether you are a financial institution or technology vendor. Add guidance about using IBM Cloud Activity Tracker Event Routing in regions where event routing is not available.

03 February 2022

- Added information on an approved variation of the VPC reference architecture to allow public internet access to the workload VPC. Added related tutorials on setting up a virtual network firewall in a new edge/transit VPC.

14 October 2021

- Revamped Best practices for software as a service. Added clarifications and detail. Also, added some new best practices derived from the Control Implementation Overview Template documents.

07 April 2021

- Initial release.

Reference architecture overview

The IBM Cloud Framework for Financial Services provides three reference architectures that can be used as a basis for meeting the security and regulatory requirements that are defined in the framework. Learn about these reference architectures in the following table.

Reference architecture	Description	
IBM Cloud® Virtual Private Cloud (VPC)	Built by using IBM Cloud® Virtual Private Cloud, which enables an enterprise to establish its own private cloud-like computing environment on shared public cloud infrastructure. The architecture has options for using one or both of IBM Cloud® Virtual Servers for Virtual Private Cloud and Red Hat® OpenShift® on IBM Cloud®	
IBM Cloud Satellite®	Built by using IBM Cloud® for VMware® Regulated Workloads	Built by using IBM Cloud for VMware Regulated Workloads, which is an extension of the VMware vCenter Server® offering. Its design extends and enhances the basic vCenter Server architecture to deliver a secure, high-performance platform.

IBM Virtual Private Cloud (VPC) architecture

VPC reference architecture for IBM Cloud for Financial Services

IBM Cloud® Virtual Private Cloud (VPC) is a public cloud offering that lets an enterprise establish its own private cloud-like computing environment on shared public cloud infrastructure. A VPC gives an enterprise the ability to define and control a virtual network that is logically isolated from all other public cloud tenants, creating a private, secure place on the public cloud. The VPC reference architecture for the IBM Cloud for Financial Services is designed to provide a framework for building a VPC-based offering according to the best practices and requirements of the IBM Cloud Framework for Financial Services. We detail this architecture and provide guidance for deploying, configuring, and managing it.

Architecture diagram

VPC reference architecture with only private access

High-level VPC reference architecture for IBM Cloud for Financial Services

Central to the architecture are two VPCs, which provide for separation of concerns between provider management functionality and consumer workloads.

Management VPC

Provides compute, storage, and network services to enable the application provider's administrators to monitor, operate, and maintain the environment.

Workload VPC

Provides compute, storage, and network services to support hosted applications and operations that deliver services to the consumer.

Other key features to note:

- Supports a single tenant.
- Resides in one or more multizone regions.
- Gives two options for compute that can be mixed and matched: IBM Cloud® Virtual Servers for Virtual Private Cloud and Red Hat® OpenShift® on IBM Cloud®.
- Enables access to the management VPC from the application provider's enterprise environment through IBM Cloud® Direct Link or IBM Cloud Virtual Private Network (VPN) for VPC.
- Provides connectivity from the consumer's enterprise environment to the workload VPC through Direct Link or VPN for VPC.
- Connects management VPC and workload VPC by using IBM Cloud® Transit Gateway.
- Allows connectivity to IBM Cloud services that use IBM Cloud Virtual Private Endpoint (VPE) for VPC.
- Encrypts data by using IBM Cloud Hyper Protect Crypto Services, which enables keep your own key (KYOK) functionality which provides technical assurance that IBM cannot access your keys.

Variation with edge or transit VPC for public internet access

The architecture in the previous section is the most secure way of enabling consumers to access the applications that are running in a workload VPC. However, there might be valid cases where it is desirable to allow consumers to access your service through the public internet. The same base architecture can be adapted to securely enable this type of access.

High-level VPC reference architecture with edge/transit VPC

The revised architecture adds:

- IBM Cloud® Internet Services (CIS) to provide global load balancing and layer 3/4 protection against distributed denial-of-service (DDoS) attacks. It also includes a web application firewall (WAF) protection and layer 7 protection against denial-of-service (DoS) attacks.
- As an alternative to CIS WAF capabilities, when advanced firewall functionalities are required, a Virtual Network Firewall software in the edge VPC can be deployed to provide web application firewall (WAF) protection and layer 7 protection against denial-of-service (DoS) attacks. The VNF will need to meet the required Financial Services controls.

See VPC architecture with virtual servers for more details on this variation.

Financial Services Validated services

Deploying the reference architecture depends upon VPC infrastructure and PaaS services that are IBM Cloud for Financial Services Validated. This means that they have evidenced compliance to the controls of the IBM Cloud Framework for Financial Services. Financial Services Validated services are designed to help address the requirements of financial institutions for regulatory compliance, security, and resiliency. When properly configured and managed, services that are Financial Services Validated work together so you can deliver a solution that conforms to the best practices of the IBM Cloud Framework for Financial Services.

⚠ **Important:** Generally speaking, you should strive to use only services which are Financial Services Validated in your solutions. However, depending on your circumstance there may be exceptions. See the best practice Use only services that are IBM Cloud for Financial Services Validated for more details and potential exceptions.

Category	Required services	Optional services
Compute [1]	• IBM Cloud® Virtual Servers for Virtual Private Cloud	• Dedicated hosts for VPC [2] • IBM Cloud Auto Scale for VPC [3]
Containers [4]	• Red Hat® OpenShift® on IBM Cloud® • IBM Cloud® Container Registry	
Networking - VPC infrastructure	• IBM Cloud® Virtual Private Cloud • IBM Cloud® Application Load Balancer for VPC • IBM Cloud® Virtual Private Network (VPN) for VPC [5] • IBM Cloud® DNS Services • IBM Cloud Virtual Private Endpoint (VPE) for VPC	
Networking - interconnectivity	• IBM Cloud® Direct Link (2.0) [6] • IBM Cloud® Transit Gateway	
Storage	• IBM® Cloud Block Storage for Virtual Private Cloud • IBM Cloud® Object Storage	
Security	• IBM Cloud® Hyper Protect Crypto Services	• IBM Cloud® Secrets Manager • IBM Cloud® App ID
Logging and monitoring	• IBM Cloud Activity Tracker Event Routing [7] • IBM Cloud Compliance Manager • IBM Cloud® Flow Logs for VPC [8] • IBM Cloud Logs	
Integration		• IBM® Event Streams for IBM Cloud®
Developer tools		• IBM Cloud® Continuous Delivery

Required and optional services for VPC reference architecture

The remainder of this section goes into more detail about how these services fit into the reference architecture.

Compute

Virtual Servers for VPC

Virtual Servers for VPC is an infrastructure-as-a-service (IaaS) offering that gives you access to all of the benefits of VPC, including network isolation, security, and flexibility. You can quickly provision instances with high network performance. When you provision an instance, you select a profile that matches the amount of memory and compute power that you need for the application that you plan to run on the instance. Instances are available on the x86 architecture.

Dedicated hosts for VPC (optional)

You can optionally use dedicated hosts for VPC. You can create a dedicated host to carve out a single-tenant compute node, free from users outside of your organization. Within that dedicated space, you can create virtual server instances according to your needs. Additionally, you can create dedicated host groups that contain dedicated hosts for a specific purpose. Because a dedicated host is a single-tenant space, only users within your account that have the required permissions can create instances on the host.

Tip: Dedicated hosts are highly recommended when you use virtual servers -- particularly for any parts of your application that process regulated data and keep it in memory.

IBM Cloud Auto Scale for VPC (optional)

With Auto Scale for VPC, you can improve performance and costs by dynamically creating virtual server instances to meet the demands of your environment. Auto Scale for VPC is highly recommended if you are using virtual servers. You set scaling policies that define your desired average utilization for metrics like CPU, memory, and network usage. The policies that you define determine when virtual server instances are added or removed from your instance group. Auto Scale for VPC is highly recommended if you are using virtual servers.

Containers

Red Hat OpenShift on IBM Cloud

Red Hat OpenShift on IBM Cloud is a managed offering to create your own Red Hat OpenShift on IBM Cloud cluster of compute hosts to deploy and manage containerized apps on IBM Cloud. Red Hat OpenShift on IBM Cloud provides intelligent scheduling, self-healing, horizontal scaling, service discovery and load balancing, automated rollouts and rollbacks, and secret and configuration management for your apps. Combined with an intuitive user experience, built-in security and isolation, and advanced tools to secure, manage, and monitor your cluster workloads, you can rapidly deliver highly available and secure containerized apps in the public cloud.

In practice, when you choose Red Hat OpenShift on IBM Cloud for your primary compute, you might also need one or more instances of Virtual Servers for VPC for other parts of the reference architecture.

IBM Cloud Container Registry

Container Registry provides a multi-tenant, highly available, scalable, and encrypted private image registry that is hosted and managed by IBM®. When you push images to Container Registry, you benefit from the built-in Vulnerability Advisor features that scan for potential security issues and vulnerabilities.

Networking - VPC infrastructure

IBM Cloud Application Load Balancer for VPC

Use Application Load Balancer for VPC (ALB) to distribute traffic among multiple server instances within the same region of your VPC. You can create a public or private ALB.

Note: Network Load Balancer for VPC (NLB) is also IBM Cloud for Financial Services Validated, but does not span zones. So, NLBs are not typically used in applications where high availability is needed.

IBM Cloud Virtual Private Network (VPN) for VPC

Use the VPN for VPC service to securely connect your VPC to another private network. Use a static, route-based VPN or a policy-based VPN to set up an IPsec site-to-site tunnel between your VPC and your on-premises private network, or another VPC.

Client VPN for VPC provides client-to-site connectivity, which allows remote devices to securely connect to the VPC network using an OpenVPN software client. This solution is useful for telecommuters who want to connect to IBM Cloud from a remote location, such as a home office, while still maintaining secure connectivity. To meet the control requirements, you should use full-tunnel mode. In full-tunnel mode, all traffic from a VPN client is routed to the

VPN server, which is more secure (especially if connecting from an untrusted network).

VPN for VPC is required to connect to the management VPC if not using Direct Link.

IBM Cloud DNS Services

DNS Services provides private DNS to VPC users. Private DNS zones are resolvable only on IBM Cloud, and only from explicitly permitted networks in an account.

IBM Cloud Virtual Private Endpoint (VPE) for VPC

With IBM Cloud Virtual Private Endpoint (VPE) for VPC you can connect to supported IBM Cloud services from your VPC network by using the IP addresses of your choosing, which is allocated from a subnet within your VPC.

VPE is an evolution of the private connectivity to IBM Cloud services. VPEs are virtual IP interfaces that are bound to an endpoint gateway created on a per service, or service instance, basis (depending on the service operation model). The endpoint gateway is a virtualized function that scales horizontally, is redundant and highly available, and spans all availability zones of your VPC. Endpoint gateways enable communications from virtual server instances within your VPC and IBM Cloud service on the private backbone. VPE for VPC gives you the experience of controlling all the private addressing within your cloud.

Networking - Interconnectivity

IBM Cloud Direct Link

Use Direct Link to seamlessly connect your on-premises resources to your cloud resources. The speed and reliability of Direct Link extends your organization's data center network and offers more consistent, higher-throughput connectivity, keeping traffic within the IBM Cloud network. Direct Link is the most secure way to enable connectivity from on-premises environments to IBM Cloud.

Direct Link is required if not using Virtual Private Network (VPN) for VPC (see next section).

IBM Cloud Transit Gateway

As the number of your VPCs grow, you need an easy way to manage the interconnection between these resources across multiple regions. Transit Gateway is designed specifically for this purpose, and is the means for connecting your management VPC to your workload VPC.

Storage

Block Storage for VPC

Block Storage for VPC provides hypervisor-mounted, high-performance data storage for your virtual server instances that you can provision within a VPC. The VPC infrastructure provides rapid scaling across zones and extra performance and security.

Block Storage for VPC is used for both primary boot volumes and secondary data volumes. Boot volumes are automatically created and attached during instance provisioning. Data volumes can be created and attached during instance provisioning as well, or as stand-alone volumes that you can later attach to an instance. To protect your data, you should use KYOK encryption with Hyper Protect Crypto Services.

IBM Cloud Object Storage

Object Storage stores encrypted and dispersed data across multiple geographic locations. Object Storage is available with three types of resiliency: Cross Region, Regional, and Single Data Center. Cross Region provides higher durability and availability than using a single region at the cost of slightly higher latency. Regional service reverses those tradeoffs, and distributes objects across multiple availability zones within a single region. If a given region or availability zone is unavailable, the object store continues to function without impediment. Single Data Center distributes objects across multiple machines within the same physical location.

Users of Object Storage refer to their binary data, such as files, images, media, archives, or even entire databases as objects. Objects are stored in a bucket, the container for their unstructured data. Buckets contain both inherent and user-defined metadata. Finally, objects are defined by a globally unique combination of the bucket name and the object key, or name.

Security

IBM Cloud Hyper Protect Crypto Services

Hyper Protect Crypto Services is a dedicated key management service and hardware security module (HSM) based on IBM Cloud. This service allows you to take the ownership of the cloud HSM to fully manage your encryption keys and to perform cryptographic operations using Keep Your Own Key (KYOK). Hyper Protect Crypto Services is also the only service in the cloud industry that is built on FIPS 140-2 Level 4-certified hardware.

IBM Cloud App ID (optional)

App ID helps developers to easily add authentication to their web and mobile apps with few lines of code, and secure their cloud-native applications and services on IBM Cloud.

Logging and monitoring

IBM Cloud Activity Tracker Event Routing

Activity Tracker Event Routing is used to collect auditable platform events that are generated by services in your IBM Cloud account. These events allow you to monitor the activity of your IBM Cloud account so that you can investigate abnormal activity and critical actions.

Activity Tracker Event Routing provides for either event routing or hosted event search. However, only the event routing features of Activity Tracker Event Routing are Financial Services Validated. In regions where it's available, you must configure Activity Tracker Event Routing to send events to Object Storage, where they must be encrypted with KYOK.

⚠ **Important:** Activity Tracker Event Routing is only available in some regions (see Locations for Activity Tracker Event Routing event routing for more details). For regions where it's not available, you must use Activity Tracker Event Routing hosted event search until Activity Tracker Event Routing is available. When event routing becomes available in those regions, you must switch to use event routing. For more information and possible exceptions, see Use only services that are IBM Cloud for Financial Services Validated.

IBM Cloud® Compliance Manager

With Compliance Manager you can embed security checks into your every day workflows to help monitor for security and compliance. By monitoring for risks, you can identify security vulnerabilities and quickly work to mitigate the impact and fix the issue. By using Compliance Manager along with external integrations (such as, OpenShift Compliance Operator (OSCO), Tanium, NeuVector, and so on), you can build a robust approach for monitoring for security and compliance issues.

IBM Cloud Flow Logs for VPC

Flow Logs for VPC enables the collection, storage, and presentation of information about the Internet Protocol (IP) traffic flowing to and from network interfaces within your VPC.

Flow Logs for VPC can help with a number of tasks, including:

- Troubleshooting why specific traffic isn't reaching an instance, which helps to diagnose restrictive security group rules
- Recording the metadata of network traffic that is reaching your instance
- Determining source and destination traffic from the network interfaces
- Adhering to compliance regulations
- Assisting with root cause analysis

IBM Cloud Logs

VPC concepts.

- Explore more detailed views of the VPC reference architecture based on compute type:
- Virtual Servers for VPC
- Red Hat OpenShift on IBM Cloud

---

1. Only required if using virtual servers instead of or in addition to Red Hat OpenShift on IBM Cloud. ↩
2. Highly recommended if using virtual servers instead of or in addition to Red Hat OpenShift on IBM Cloud. ↩
3. Highly recommended if using virtual servers instead of or in addition to Red Hat OpenShift on IBM Cloud. ↩
4. Only required if using containers instead of or in addition to virtual servers. ↩
5. Only choose one of Direct Link and VPN for VPC. ↩
6. Only choose one of Direct Link and VPN for VPC. ↩
7. Only the event routing features of Activity Tracker have been Financial Services Validated. ↩
8. Only required if using virtual servers instead of or in addition to Red Hat OpenShift on IBM Cloud. ↩
9. Only required if using containers instead of or in addition to virtual servers. ↩
10. Installing your own software is recommended when there is not yet an IBM Cloud service that is Financial Services Validated. However, when

installing your own software you are still responsible for the controls of the IBM Cloud Framework for Financial Services if you are seeking your own Financial Services Validated designation. ↩

11. Only the event routing features of Activity Tracker have been Financial Services Validated. ↩

VPC concepts

Now that you've seen the high-level VPC reference architecture for IBM Cloud for Financial Services, there are a number of important VPC concepts that will help you understand the components within the management and workload VPCs as we dive into lower-level details. Each VPC is a virtual network that is linked to your customer account. It gives you cloud security, with the ability to scale dynamically, by providing fine-grained control over your virtual infrastructure and your network traffic segmentation.

Network

Regions, zones, and subnets

All VPCs exist in a single region. A *region* is an abstraction that is related to the geographic area in which a VPC is deployed. Each multizone region contains multiple zones. A *zone* is another an abstraction that refers to the physical data center that hosts the compute, network, and storage resources, as well as the related cooling and power, which provides services and applications.

VPCs that are created in a one of the multizone regions can span multiple zones to facilitate high availability. For IBM Cloud for Financial Services, it is recommended to use at least three zones in each VPC to help ensure resiliency.

Subnets consist of a specified IP address range (CIDR block). Subnets are bound to a single zone, and they cannot span multiple zones or regions. Subnets in the same VPC are connected to each other.

For more information, see About networking.

Access control lists and security groups

Access control lists (ACLs) and security groups provide ways to control the traffic across the subnets and instances in your VPC by using rules that you specify. ACLs control traffic to and from a subnet, while security groups control the traffic at the virtual server instance (VSI) level.

For more information, see Security in your VPC.

Compute

Virtual Servers for VPC

Virtual Servers for VPC gives you access to all of the benefits of IBM Cloud VPC, including network isolation, security, and flexibility. When you provision an instance, you select a profile that matches the amount of memory and compute power that you need for the application that you plan to run on the instance. Instances are available on the x86 architecture.

See About virtual server instances for VPC for more details.

Dedicated hosts

Provision your virtual servers on a dedicated host. Dedicated hosts are used to carve out a single-tenant compute node, free from users outside of your organization. Within that dedicated space, you can create virtual server instances according to your needs. Additionally, you can create dedicated host groups that contain dedicated hosts for a specific purpose. Because a dedicated host is a single-tenant space, only users within your account that have the required permissions can create instances on the host.

See Creating dedicated hosts and groups for more details.

Red Hat OpenShift on IBM Cloud

You can also use Red Hat OpenShift on IBM Cloud to run applications. It is also backed by virtual server instances, but, in that case, the virtual service instances are managed by IBM.

See VPC reference architecture with Red Hat OpenShift on IBM Cloud for more details.

Storage

Block Storage for VPC provides hypervisor-mounted, high-performance data storage for your virtual server instances that you can provision within a VPC. The VPC infrastructure provides rapid scaling across zones and extra performance and security.

Block Storage for VPC is used for both primary boot volumes and secondary data volumes. Boot volumes are automatically created and attached during

instance provisioning. Data volumes can be created and attached during instance provisioning as well, or as stand-alone volumes that you can later attach to an instance. To protect your data, you should use KYOK encryption with Hyper Protect Crypto Services.

See About Block Storage for VPC for more details.

Next steps

- Take a deeper look at the VPC reference architecture with virtual servers.

Single-region IBM Cloud for Financial Services reference architecture for VPC with Virtual Servers for VPC

Now that you've seen the high-level VPC reference architecture for IBM Cloud for Financial Services and looked at some VPC concepts, we will take a deeper look at some of the details of a reference architecture that uses virtual servers.

This architecture shows a deployment of the VPC that uses Virtual Servers for VPC as the primary compute.

Architecture diagram

Single-region IBM Cloud for Financial Services reference architecture for VPC with Virtual Servers for VPC

Management VPC

Regions and zones

The management VPC provides compute, storage, and network services to enable application provider administrators to monitor, operate, and maintain the environment. The intent is to completely isolate management operations from the VPC running consumer workloads.

The management VPC is distributed across three zones in one multizone region (MZR).

Subnets for management tools

There could be multiple subnets with different ACLs and multiple security groups in the VPC.

The top subnet in each zone contains an arbitrary number of virtual server instances that use Block Storage for VPC. Security group is used to control the access to these instances, and they are where your management tools run.

The lower subnets in the diagram relate to inbound connectivity, and we'll cover more about that in the next sections.

Connectivity inbound to VPC

Connectivity from your application provider's enterprise environment to the management VPC is accomplished by using IBM Cloud Direct Link or IBM Cloud Virtual Private Network (VPN) for VPC. You must use one or the other.

Direct Link is the most secure way to enable connectivity from the application provider's on-premises environment to the management VPC. The speed and reliability of Direct Link extends your organization’s data center network and offers more consistent, higher-throughput connectivity, keeping traffic within the IBM Cloud network. When using Direct Link, a private Application Load Balancer for VPC (ALB) is used to distribute traffic among multiple server instances within the same region of your VPC.

An alternative connectivity pattern requires use of the VPN for VPC service to securely connect from your private network to the management VPC. VPN for VPC can be used as a static, route-based VPN or a policy-based VPN to set up an IPsec site-to-site tunnel between your VPC and your on-premises private network, or another VPC. When using VPN for VPC, you need to place the gateway in a subnet (shown in the lower left subnet in the diagram).

Client VPN for VPC provides client-to-site connectivity, which allows remote devices to securely connect to the VPC network using an OpenVPN software client. This solution is useful for telecommuters who want to connect to IBM Cloud from a remote location, such as a home office, while still maintaining secure connectivity. To meet the control requirements, you should use full-tunnel mode. In full-tunnel mode, all traffic from a VPN client is routed to the VPN server, which is more secure (especially if connecting from an untrusted network).

Bastion host

Regardless of whether you are using Direct Link or VPN for VPC to connect to the management VPC, you need to ensure that all traffic is routed through a bastion host with session recording. The bastion host solution is depicted in the two rightmost lower subnets in the diagram.

Connectivity between VPCs

The management VPC needs to connect to the workload VPC to deploy, configure, and operate the components and workloads that are found in the

workload VPC. Transit Gateway is designed specifically for this purpose, and is the means for connecting your management VPC to your workload VPC.

Workload VPC

The workload VPC provides compute, storage, and network services to support hosted applications and operations that deliver services to the consumer. Let's take a closer look at the components within the VPC.

Regions and zones

Just like the management VPC, the workload VPC is spread across three zones. The workload VPC should be created in the same MZR as the management VPC.

Connectivity to workload VPC

If the consumer is in the same organization as the application provider, then just like for the management VPC, Direct Link can provide access to the workload VPC. Alternatively, VPN for VPC can be used for site-to-site or client-to-site VPN connectivity.

Storage and encryption

Block Storage for VPC

Block Storage for VPC is automatically created for the primary book volume when you create a virtual server. You can also create more data volumes by using Block Storage for VPC within any subnet.

By default, all boot and data volumes are encrypted at rest with IBM-managed encryption. There is no additional cost for this service. IBM-managed encryption uses the following industry standard protocols:

- AES-256 encryption
- Keys are managed in-house with Key Management Interoperability Protocol (KMIP)

For all sensitive and consumer-owned data, the application provider should use customer-managed encryption. With this encryption method you can bring your own customer root key (CRK) to the cloud or have a key management service (KMS) generate a key for you. For IBM Cloud for Financial Services, the supported key management service is Hyper Protect Crypto Services (HPCS). Root keys encrypt volume and custom image passphrases with envelope encryption, a process that encrypts a key with another key.

Object Storage

Object Storage is an alternative storage option that is useful for certain use cases, including backup and recovery, data archiving, cloud-native application building, and AI and big data analytics. Object Storage stores encrypted and dispersed data across multiple geographic locations.

By default, all objects that are stored in Object Storage are encrypted by using randomly generated keys and an all-or-nothing-transform (AONT). While this default encryption model provides at-rest security, financial service workloads need full control over the data encryption keys used. Again, Hyper Protect Crypto Services should be used for this purpose.

Using IBM Cloud services outside of a VPC

With IBM Cloud Virtual Private Endpoint (VPE) for VPC you can connect to supported IBM Cloud services from your VPC network by using the IP addresses of your choosing, which is allocated from a subnet within your VPC. In the reference architecture diagram, VPEs appear in the middle subnets of the workload VPC.

VPE is an evolution of the private connectivity to IBM Cloud services. VPEs are virtual IP interfaces that are bound to an endpoint gateway created on a per service, or service instance, basis (depending on the service operation model). The endpoint gateway is a virtualized function that scales horizontally, is redundant and highly available, and spans all availability zones of your VPC. Endpoint gateways enable communications from virtual server instances within your VPC and IBM Cloud service on the private backbone. VPE for VPC gives you the experience of controlling all the private addressing within your cloud.

Variation with edge/transit VPC for public internet access

You might want to allow consumers to access your service through the public internet. This base architecture can be adapted to securely enable this type of access as shown in the following diagram, which introduces a new edge VPC. The request from the consumer gets routed through Cloud Internet Service's global load balancer, through a public load balancer in the edge VPC, and then to the private load balancer within the workload VPC. IBM Cloud Internet Services provides Domain Name Service (DNS), Global Load Balancer (GLB), DDoS protection, Web Application Firewall (WAF), Transport Layer Security (TLS), Rate Limiting, Smart Routing, and Caching, and can be used for public internet traffic.

Single-region IBM Cloud for Financial Services reference architecture for VPC with BIG-IP

For complete details on this variation of the architecture, see Consumer connectivity to workload VPC.

Next steps

If you plan to use Red Hat OpenShift on IBM Cloud, explore a more detailed view of the VPC reference architecture with Red Hat OpenShift on IBM Cloud

If you don't plan to use Red Hat OpenShift on IBM Cloud, you can skip ahead to learn more about deployment in Setup environment for deployment and configuration.

Virtual Servers for VPC and Power Virtual Server reference architecture for IBM Cloud for Financial Services

This solution pattern contains the design and architecture decisions for cloud native and Power Virtual Server workloads in IBM Cloud for Financial Services.

IBM Power is a family of high-performance servers that are designed for running large-scale data-driven and mission-critical workloads and are known for their scalability, reliability, sustainability, and performance. IBM® Power® Virtual Server is a Power Systems offering in IBM Cloud. As stated in the documentation, Power Virtual Server are located in the IBM data centers, distinct from the IBM Cloud servers with separate networks and direct-attached storage. The internal networks are fenced but offer connectivity options to IBM Cloud infrastructure or on-premises environments. This infrastructure design enables Power Virtual Server to maintain key enterprise software certification and support as the Power Virtual Server architecture is identical to certified on-premises infrastructure. This documentation page discusses benefits, use cases, workloads and other information on Power Virtual Server.

The IBM Cloud Framework for Financial Services provides different flavors of reference architectures to be used as the basis for meeting the security and regulatory requirements defined in the framework. This document extends the Virtual Servers for VPC based reference architecture and provides a solution design for deploying sensitive enterprise workloads on Power Virtual Server in cloud, especially for financial and regulated industries, by adopting the IBM Cloud for Financial Services Framework and follow the principals and best practices of the framework.

Architecture diagram

The diagram below represents the architecture for secure Power Virtual Server workloads in IBM Cloud and is an extension of the Virtual Servers for VPC reference architecture for IBM Cloud for Financial Services.

IBM Cloud for Financial Services reference architecture for VPC and Power Virtual Server

Central to the architecture are three VPCs, which provide segmentation for edge traffic control, management functionality, and consumer workloads.

Management VPC

> Provides compute, storage, and network services to enable the client or service provider's administrators to monitor, operate, and maintain the environment.

Workload VPC

> Provides compute, storage, and network services to support hosted applications and operations that deliver services to the consumer.

Edge VPC

> The edge VPC is used to enhance boundary protection for the workload VPC, by allow consumers to access public facing interface through internet.

This reference architecture is an extension of the VPC reference architecture for IBM Cloud for Financial Services. Here are some key features:

- Supports a single tenant.
- Resides in one or more multizone regions. When resources are deployed in multiple regions for resiliency purpose, Global Transit Gateway can be used to connect traffic across regions.
- Enables access to the management VPC from the application provider's enterprise environment through IBM Cloud® Direct Link or Virtual Private Network (VPN) for VPC.
- Manages traffic flow from outside IBM Cloud via Edge VPC.
- Supports multiple subnets, ACLs and Security Groups to manage traffic flow to the subnets and to the instances. Load balancers can be used.
- Supports Bastion host to manage access to internal servers. Power Virtual Server instances should not be exposed directly to external world.
- Allows connectivity to IBM Cloud services that use Virtual Private Endpoint (VPE) for VPC. Custom DNS server and proxy server can be used.
- IBM Cloud® Internet Services to provide global load balancing and layer 3/4 protection against distributed denial-of-service (DDoS) attacks.
- Virtual network firewall software in the Edge VPC to provide web application firewall (WAF) protection and layer 7 protection against denial-of-

service (DoS) attacks.

- Power Virtual Server are located in the IBM data centers, distinct from the IBM Cloud servers with separate networks and direct-attached storage. Power Edge Router (PER) (or Cloud Connections where PER is not available) provides connection between Power Virtual Servers and IBM Cloud resources.
- Power workloads can be deployed on Power Virtual Server instances, cloud native workloads can be deployed in VPC, and Transit Gateway can be used to define and control communication between resources on the IBM Cloud network. Separates enterprise traffic and IBM Cloud internal traffic via different Transit Gateway.
- Power workloads can be deployed on Power Virtual Server instances in different regions. Global Transit Gateway can be used in this case.
- Hyper Protect Crypto Services or Key Protect can be integrated with Power Virtual Server.
- IBM Cloud for Financial Services Validated services are used. Compliance Manager provides security and compliance support.

Design concepts

Following the Architecture Framework, this document covers the following solution aspects and domains:

Architecture design scope

Requirements

The following table outlines the requirements that are addressed in this architecture.

Aspect	Requirements
Compute	Provide properly isolated compute resources with adequate compute capacity for the applications.
Storage	Provide storage that meets the application and database performance requirements.
Networking	Deploy workloads in isolated environment and enforce information flow policies. Provide secure, encrypted connectivity to the cloud's private network for management purposes. Distribute incoming application requests across available compute resources.
Security	Ensure all operator actions are executed securely through a bastion host. Protect the boundaries of the application against denial-of-service and application-layer attacks. Encrypt all application data in transit and at rest to protect from unauthorized disclosure. Encrypt all security data (operational and audit logs) to protect from unauthorized disclosure. Encrypt all data using customer managed keys to meet regulatory compliance requirements for additional security and customer control. Protect secrets through their entire lifecycle and secure them using access control measures. Firewalls must be restrictively configured to prevent all traffic, both inbound and outbound, except that which is required, documented, and approved.
DevOps	Delivering software and services at the speed the market demands requires teams to iterate and experiment rapidly. They must deploy new versions frequently, driven by feedback and data.
Resiliency	Support application availability targets and business continuity policies. Ensure availability of the application in the event of planned and unplanned outages. Backup application data to enable recovery in the event of unplanned outages. Provide highly available storage for security data (logs) and backup data.
Service Management	Monitor system and application health metrics and logs to detect issues that might impact the availability of the application. Generate alerts/notifications about issues that might impact the availability of applications to trigger appropriate responses to minimize down time. Monitor audit logs to track changes and detect potential security problems. Provide a mechanism to identify and send notifications about issues found in audit logs.

Requirements

Components

The following table outlines the products or services used in the architecture for each aspect.

Aspects	Architecture components	How the component is used

Application platforms	Red Hat® OpenShift® on IBM Cloud®	Deploy and secure enterprise workloads
	IBM Cloud® Code Engine	Run your application, batch job, or container on a managed serverless platform
	IBM Cloud® Container Registry	Securely store container images and monitor their vulnerabilities in a private registry
Data	IBM® Event Streams for IBM Cloud®	Leverage fully managed Kafka service to build intelligent applications that react to events in real time
Compute	IBM Cloud® Virtual Servers for Virtual Private Cloud	Virtual machines with faster provisioning, higher performance, and enhanced isolation
	Dedicated hosts for VPC	Provision single-tenant hosts that offer dedicated resources and maximum control over instance placement
	IBM® Power® Virtual Server	Virtual Server offering on IBM Power systems
Storage	IBM Cloud® Object Storage	Provide flexible, cost-effective, and scalable cloud storage for unstructured data
	IBM® Cloud Block Storage for Virtual Private Cloud	Persistent storage for use as boot and data storage for Virtual Servers in a VPC network
Networking	IBM Cloud® Virtual Private Cloud	Fully customizable, software-defined virtual network with superior isolation
	Application Load Balancer for VPC	Distribute layer 7 and 4 traffic among server instances within the same region of your VPC and support Secure Sockets Layer (SSL) offloading
	Network Load Balancer for VPC	Distribute layer 4 traffic among multiple server instances within the same region of your VPC
	Virtual Private Network (VPN) for VPC	Connect your on-premises network to the IBM Cloud VPC network
	Client VPN for VPC	Securely connect to your IBM Cloud resources from anywhere through a client-to-site VPN server
	IBM Cloud® Internet Services	Offers capabilities to enhance your workflow
	IBM Cloud® DNS Services	Manage hostnames and IP addresses while limiting access to the DNS records from permitted networks only
	Virtual Private Endpoint (VPE) for VPC	Delivers private connectivity through Endpoint Gateways to IBM Cloud Services utilizing client assigned IP addresses from within the VPC
	IBM Cloud® Direct Link	Establish and deliver connectivity to IBM Cloud
	IBM Cloud® Transit Gateway	Creates secure connectivity between your networks within IBM Cloud
Security	IBM Cloud® App ID	User Authentication and User Profiles for your apps
	Hyper Protect Crypto Services	Keep Your Own Key for cloud data encryption with a dedicated key management service built on FIPS 140-2 Level 4 certified HSM
	IBM® Key Protect	Create or manage cryptographic keys in the cloud or in Satellite to protect data at rest

	IBM Cloud® Secrets Manager	Create, lease, and centrally manage secrets that are used in your apps and services
DevOps	IBM Cloud® Continuous Delivery	Support DevOps best practices by using Git, issue tracking, source code vulnerability analysis, and CI/CD pipelines in the Cloud
	Toolchain	Automates the tasks of developing and deploying your app
Resiliency	IBM Cloud Backup for VPC	Provides the ability to schedule VPC block storage snapshot backups and manage retention through backup policies
	Block Storage Snapshots for VPC	Back up block storage volumes to IBM Cloud Object Storage with this regional snapshot service
Service management	IBM Cloud Activity Tracker Event Routing	Collect auditable platform events that are generated by services in your IBM Cloud account
	IBM Cloud® Compliance Manager	Manage your security and compliance posture
	Flow Logs for VPC	enable the collection, storage, and presentation of information about the Internet Protocol (IP) traffic going to and from network interfaces within your Virtual Private Cloud (VPC)

Components

Single-region IBM Cloud for Financial Services reference architecture for VPC with Red Hat OpenShift on IBM Cloud

If you want to use containers, you can add Red Hat OpenShift on IBM Cloud to your VPC. Except for the addition of Red Hat OpenShift on IBM Cloud, you use the same architectural patterns and components that were described for the Single-region IBM Cloud for Financial Services reference architecture for VPC with Virtual Servers for VPC.

The following diagram shows a more detailed view of both the management and workload VPCs when Red Hat OpenShift on IBM Cloud is introduced.

Architecture diagram

Single region IBM Cloud for Financial Services reference architecture for VPC with Red Hat OpenShift on IBM Cloud

You can choose to use Red Hat OpenShift on IBM Cloud alongside (or instead of) virtual servers in either or both VPCs. Even though it is shown in the diagram as an option, it is not required to put Red Hat OpenShift on IBM Cloud in your management VPC.

Red Hat OpenShift on IBM Cloud concepts

Red Hat OpenShift on IBM Cloud is a managed offering to create your own Red Hat OpenShift on IBM Cloud cluster of compute hosts to deploy and manage containerized apps on IBM Cloud. Red Hat OpenShift on IBM Cloud provides intelligent scheduling, self-healing, horizontal scaling, service discovery and load balancing, automated rollouts and rollbacks, and secret and configuration management for your apps. Combined with an intuitive user experience, built-in security and isolation, and advanced tools to secure, manage, and monitor your cluster workloads, you can rapidly deliver highly available and secure containerized apps in the public cloud.

Red Hat OpenShift on IBM Cloud extends the Kubernetes platform with built-in software to enhance app lifecycle development, operations, and security. The master nodes are entirely IBM's responsibility, while there is shared responsibility for the worker nodes.

⚠ **Important:** For IBM Cloud for Financial Services, you should provision Red Hat OpenShift on IBM Cloud only in a VPC and not in classic infrastructure.

For more information, see Understanding Red Hat OpenShift on IBM Cloud.

Next steps

- Setup environment for deployment and configuration.

IBM Cloud Satellite architecture

Satellite reference architecture for IBM Cloud for Financial Services

With shared responsibility model to fulfill the best practices and requirements of the IBM Cloud Framework for Financial Services. This document describes the architecture and provides guidance for deploying, configuring, and managing it.

Satellite achieves this distributed cloud architecture by providing an API-based suite of tools that you can use to represent your on-premises data center, a public cloud provider, or an edge network as a Satellite location. You fill the Satellite location with your own host machines that meet the minimum host requirements. Then, these hosts provide the compute power to run IBM Cloud services, such as workloads in managed Red Hat OpenShift clusters.

Solution design

When combined with other IBM Cloud services that run in the cloud and other applications in your on-premises environment, the hybrid architecture of Satellite can be used to build robust solutions that are appropriate for financial and other regulated workloads. This hybrid architecture is shown in the following diagram.

High-level Satellite reference architecture for IBM Cloud for Financial Services

The following table provides an index to the components in the diagram.

IBM Cloud	Satellite location on-premises	Supporting components on-premises
1. Satellite link tunnel server	5. Storage	11. Edge plane
2. Monitoring	6. Satellite hosts	12. Management plane
3. Satellite Config	7. Link connector	13. Other application services
4. IBM Cloud services	8. Monitoring	
	9. Cluster 1 master 9. Cluster 2 master	
	10. Workload cluster 1 10. Workload cluster 2	

Index of components in the Satellite reference architecture diagram

Note: Satellite also supports deployments to other public cloud providers. However, we do not explore that feature in this reference architecture.

Financial Services Validated services

The reference architecture depends upon services that are IBM Cloud for Financial Services Validated. These services are designated to have evidenced compliance to the controls of the IBM Cloud Framework for Financial Services. Financial Services Validated services are designed to help address the requirements of financial institutions for regulatory compliance, security, and resiliency.

When properly configured and managed, services that are Financial Services Validated work together so you can deliver a solution that conforms to the best practices of the IBM Cloud Framework for Financial Services. The following table shows the required services that are included in the Satellite reference architecture and the complementary optional services that are available.

Important: Generally speaking, you should strive to use only services which are Financial Services Validated in your solutions. However, depending on your circumstance there may be exceptions. See the best practice Use only services that are IBM Cloud for Financial Services Validated for more details and potential exceptions.

Category	Required services	Optional services
Satellite	• IBM Cloud Satellite®	

Containers	• Red Hat® OpenShift® on IBM Cloud® [1] • IBM Cloud® Container Registry	
Storage	• IBM Cloud® Object Storage	
Security	• IBM Cloud Hyper Protect Crypto Services	• IBM Cloud® Secrets Manager • IBM Cloud® App ID
Logging and monitoring	• IBM Cloud® Activity Tracker Event Routing [2] • IBM Cloud® Compliance Manager	
Integration		• IBM® Event Streams for IBM Cloud®
Developer tools		• IBM Cloud® Continuous Delivery

Required and optional services for Satellite reference architecture

The remainder of the topic goes into more detail about how these services fit into the reference architecture.

Core

IBM Cloud Satellite

Most of the services that are mentioned run entirely in IBM Cloud but two exceptions:

- Satellite, which as described earlier, has some components that run in IBM Cloud and some components that run in your Satellite location.
- Red Hat OpenShift on IBM Cloud, which refers to the Red Hat OpenShift on IBM Cloud-enabled type of the service, runs in your Satellite location.

These exceptions have implications for what it means to be Financial Services Validated and that is represented in the following diagram and details.

Financial Services Validated components in the Satellite reference architecture

- Satellite components that run in IBM Cloud are validated to the same IBM Cloud Framework for Financial Services controls as any other Financial Services Validated product.
- Satellite components and IBM Cloud services that run in your on-premises environment that do not have a dependency on your configuration or underlying internal network and infrastructure components, are also validated to the applicable controls that were validated for these components in IBM Cloud.
- On-premises infrastructure and your on-premises workloads are your responsibility and are not validated by IBM. You are responsible for managing infrastructure and physical controls and controls that are related to deployment, configuration, and management of your workloads.

Note: Several Satellite-enabled services can run in your Satellite location. However, only Red Hat OpenShift on IBM Cloud is Financial Services Validated. So, the others are not yet part of this reference architecture.

Containers

Red Hat OpenShift on IBM Cloud

You use Red Hat OpenShift on IBM Cloud to run your application workloads in your Satellite location. Red Hat OpenShift on IBM Cloud is a managed offering to create your own Red Hat OpenShift on IBM Cloud cluster of compute hosts to deploy and manage containerized apps on IBM Cloud. Red Hat OpenShift on IBM Cloud provides intelligent scheduling, self-healing, horizontal scaling, service discovery and load balancing, automated rollouts and rollbacks, and secret and configuration management for your apps. Combined with an intuitive user experience, built-in security and isolation, and advanced tools to secure, manage, and monitor your cluster workloads, you can rapidly deliver highly available and secure containerized apps in the public cloud.

IBM Cloud Container Registry

Container Registry provides a multi-tenant, highly available, scalable, and encrypted private image registry that is hosted and managed by IBM®. When you push images to Container Registry, you benefit from the built-in Vulnerability Advisor features that scan for potential security issues and vulnerabilities.

Storage

IBM Cloud Object Storage

Object Storage stores encrypted and dispersed data across multiple geographic locations. Object Storage is available with three types of resiliency: Cross Region, Regional, and Single Data Center. Cross Region provides higher durability and availability than using a single region at the cost of slightly higher latency. Regional service reverses those tradeoffs, and distributes objects across multiple availability zones within a single region. If a given region or availability zone is unavailable, the object store continues to function without impediment. Single Data Center distributes objects across multiple machines within the same physical location.

Users of Object Storage refer to their binary data, such as files, images, media, archives, or even entire databases as objects. Objects are stored in a bucket, the container for their unstructured data. Buckets contain both inherent and user-defined metadata. Finally, objects are defined by a globally unique combination of the bucket name and the object key, or name.

Security

IBM Cloud Hyper Protect Crypto Services

Hyper Protect Crypto Services is a dedicated key management service and hardware security module (HSM) based on IBM Cloud. This service allows you to take the ownership of the cloud HSM to fully manage your encryption keys and to perform cryptographic operations using Keep Your Own Key (KYOK). Hyper Protect Crypto Services is also the only service in the cloud industry that is built on FIPS 140-2 Level 4-certified hardware.

IBM Cloud App ID (optional)

App ID helps developers to easily add authentication to their web and mobile apps with few lines of code, and secure their cloud-native applications and services on IBM Cloud.

Logging and monitoring

IBM Cloud Activity Tracker Event Routing

Activity Tracker Event Routing is used to collect auditable platform events that are generated by services in your IBM Cloud account. These events allow you to monitor the activity of your IBM Cloud account so that you can investigate abnormal activity and critical actions.

Activity Tracker Event Routing provides for either event routing or hosted event search. However, only the event routing features of Activity Tracker Event Routing are Financial Services Validated. In regions where it's available, you must configure Activity Tracker Event Routing to send events to Object Storage, where they must be encrypted with KYOK.

⚠ **Important:** Activity Tracker Event Routing is only available in some regions (see Locations for Activity Tracker Event Routing event routing for more details). For regions where it's not available, you must use Activity Tracker Event Routing hosted event search until Activity Tracker Event Routing is available. When event routing becomes available in those regions, you must switch to use event routing. For more information and possible exceptions, see Use only services that are IBM Cloud for Financial Services Validated.

IBM Cloud Compliance Manager

With Compliance Manager you can embed security checks into your every day workflows to help monitor for security and compliance. By monitoring for risks, you can identify security vulnerabilities and quickly work to mitigate the impact and fix the issue. By using Compliance Manager along with external integrations (such as, OpenShift Compliance Operator (OSCO), Tanium, NeuVector, and so on), you can build a robust approach for monitoring for security and compliance issues.

Integration

IBM Event Streams for IBM Cloud (optional)

Event Streams is a high-throughput message bus built with Apache Kafka. It is optimized for event ingestion into IBM Cloud and event stream distribution between your services and applications.

You can use Event Streams to complete the following tasks:

- Offload work to back-end worker applications.
- Connect event streams to streaming analytics to realize powerful insights.

- Publish event data to multiple applications to react in real time.

Satellite components in IBM Cloud

When you provision the Satellite service, the Satellite control plane master is created in your account. The following table describes the main components of the control plane master.

Master components	Description
Satellite Link tunnel	The Satellite Link tunnel server creates a secure TLS connection to the Satellite Link connector that runs on the control plane worker nodes in your Satellite location. Three tunnels are created between the tunnel server and the connector to support redundancy across the three availability zones of your location. All communication that leaves and enters your location is proxied by the Link tunnel server, and the connection metadata that is captured and monitored by IBM to detect malicious activity. To control the network connections between the workloads that run in your location and the IBM Cloud multizone metro that manages your location, you can set up Satellite Link endpoints. For more information, see Connecting your Satellite location and IBM Cloud with endpoints.
IBM Cloud Monitoring	The IBM Cloud Monitoring component monitors the compute capacity in your Satellite location and the components that run in your Satellite control plane to detect issues and automatically resolve them if possible. These actions can include assigning more hosts to the control plane or restart components that keep on failing. For issues that cannot be resolved, IBM Site Reliability Engineers are automatically informed for further investigation.
Satellite Config	Based on the Razee open source project, Satellite Config is a continuous delivery tool that you can use to consistently roll out versions of your apps across Red Hat OpenShift on IBM Cloud clusters in Satellite locations. For more information, see Deploying Red Hat OpenShift on IBM Cloud resources across clusters with Satellite configurations.

Overview of Satellite control plane master components

Satellite location components

A location is a representation of an environment in your on-premises data center that you want to bring IBM Cloud services to so that you can run workloads in your own environment. You create the location based on at least three separate zones of your backing infrastructure environment, and attach host machines from across these zones in your infrastructure to the location. The Satellite console provides a single pane of glass to manage the workloads that run across the infrastructure in your locations. For more information, see Planning your infrastructure environment for Satellite and Setting up Satellite locations.

Infrastructure

Satellite location-underlying infrastructure

You are responsible for providing the underlying infrastructure components in your on-premises data center that the Satellite location uses. These components include:

- Virtual and bare metal servers
- Virtual storage
- Virtual network
- Hypervisor
- Physical servers and memory
- Physical storage
- Physical network and devices
- Facilities and data centers

Satellite infrastructure components

The following table describes Satellite infrastructure components that run on top of the underlying infrastructure that you provide.

Infrastructure components	Description

Hosts	Hosts are machines that reside in your infrastructure provider, across at least three separate zones, and must meet the minimum host requirements. After attaching the hosts to a Satellite location, you assign the hosts to the Satellite clusters to provide the computing power to run your application workloads. For more information, see Setting up Satellite hosts.
Storage	Satellite storage uses Satellite Config to provide a convenient way to install various storage drivers in Red Hat OpenShift on IBM Cloud clusters across your Satellite locations, by using storage templates. The storage templates are provided and tested by the vendors. After you install Satellite storage, your cluster users can use Kubernetes persistent volume claims (PVCs) to order and save their application data in persistent storage. For more information, see Understanding Satellite storage templates.

Overview of the Satellite location infrastructure components

Satellite control plane worker nodes

The following table describes the components that run as Satellite control plane worker nodes.

Worker node component	Description
Satellite Link connector	The Satellite link connector component is the main gateway for any communication between your Satellite location and IBM Cloud. All workloads that run in your location and that must connect to an app, service, or server that runs in IBM Cloud must send their requests to the Satellite Link connector. The Satellite Link connector securely forwards your request to the Satellite Link tunnel server where the request is proxied and forwarded to the destination target. To enable DNS resolution between your Satellite location and IBM Cloud, you must create a Satellite Link endpoint. For more information, see Connecting your Satellite location and IBM Cloud with endpoints. Workload cluster communication that leaves and enters your location is proxied by the Link tunnel server, and network traffic on this connection can be monitored and audited. The host requirements guide details any current exceptions to this policy.
IBM Cloud Monitoring	The IBM Cloud Monitoring component monitors the compute capacity in your Satellite location and the components that run in your Satellite control plane to detect issues and automatically resolve them if possible. These actions can include assigning more hosts to the control plane or restart components that keep on failing. For issues that cannot be resolved, IBM Site Reliability Engineers are automatically informed for further investigation.
Cluster master	When you create a Satellite cluster in your location, the master of this cluster is deployed onto your Satellite control plane worker nodes to allow communication to IBM Cloud and monitoring through IBM. For more information, see Creating Satellite clusters.

Overview of Satellite control plane worker node components

Workload clusters

You will create Red Hat OpenShift on IBM Cloud clusters in your Satellite location to run your application workloads. You will use the hosts of your own infrastructure that you added to your location as the worker nodes for the cluster. See Creating Red Hat OpenShift clusters in Satellite.

Other supporting components on-premises

In addition to your Satellite locations, you will need other components that run on-premises to complete the reference architecture including an edge plane, management plane, and (optionally) other application services. One of the most important features of the reference architecture is the definition of physical and logical separation between the edge plane, management plane, and Satellite workload clusters. This is used to ensure separation of application workloads from management functions and to isolate security functions from non-security functions.

The following table describes these components in more detail.

Component	Description
Edge plane	• Improves network isolation and boundary protection for your workloads. • Includes tools like virtual network firewalls, and load balancers.
Management plane	• Provides separation between your workloads and management operations. • Includes a bastion host and tools like CI/CD, logging, and monitoring.

Other application services (optional)	• Provides any other on-premises services that are needed by your workloads, such as databases and APIs.

Overview of other on-premises components outside of Satellite location

Next steps

- Satellite use cases
- Satellite shared responsibility model
- Satellite architecture best practices

---

1. Satellite-enabled service which runs in your Satellite location. ↩
2. Only the event routing features of Activity Tracker have been Financial Services Validated. ↩

Satellite reference architecture use cases

Use cases for the Satellite reference architecture for the IBM Cloud for Financial Services typically fall into the four categories in the following table.

Use Case	Description
Platform modernization	Innovate faster and free your team to develop applications with Red Hat OpenShift on IBM Cloud on-premises. Establish container portability and PaaS foundation for developer services and digital supply chain operations. Achieve improved cybersecurity with industry-leading Keep Your Own Key encryption.
Concentration risk	Mitigate regulatory concerns about concentration risk by using a hybrid cloud deployment model that uses Satellite to extend secure Financial Services Validated offerings from IBM Cloud to your data center.
Data residency	Manage local data requirements — storage (residency) and legal protections (sovereignty) — by keeping data in-country and accessing public cloud services. Improve data privacy (cybersecurity) for regulated data with industry-leading Keep Your Own Key encryption.
AI and automation	Gain new insights and process automation by enhancing databases on-premises with IBM Cloud Pak AI and automation services that are deployed with Red Hat OpenShift on IBM Cloud. Maintain control of data on-premises for performance, sovereignty, and regulatory requirements.

Overview of use cases for the Satellite reference architecture

Two uses cases from real financial institutions are presented in the following sections.

US regional bank

During a time-sensitive merger, a US regional bank needed to integrate two entities' IT operations securely and bring consistency across container platforms in on-premises branch locations without adding overhead. Satellite enabled the customer to establish a location that used its existing data center infrastructure and create multiple Red Hat OpenShift on IBM Cloud clusters so that teams could build and host containerized workloads in a consistent way and pay only for what they consumed. This scenario is an example of using Satellite for platform modernization.

The following diagram shows a high-level overview of the architecture that is used by this financial institution.

High-level modification of Satellite reference architecture for the first institution

Key features of their approach include:

- Using VMWare or other virtualization infrastructure to provide the hosts in the Satellite location.
- Leveraging a proxy that they manage for traffic to and from the Satellite location.
- Running Satellite-enabled Red Hat OpenShift on IBM Cloud for their workloads and exploring other Satellite-enabled services in later phases when they become Financial Services Validated.
- Hosting on-premises databases on virtual machines outside the Satellite location within the enterprise.
- Placing management tools in virtual servers within their IBM Cloud account.
- Using the following services on IBM Cloud:
  - IAM for identity and access management

- Activity Tracker Event Routing to collect IBM Cloud audit events
- Object Storage for backup and archiving data

Multinational bank

A multinational bank wanted to ensure they were following local requirements for data residency by keeping data in country when they were accessing public cloud services. The Satellite reference architecture was a perfect fit for them.

The following diagram shows a high-level overview of the architecture that is used by this second institution.

High-level modification of Satellite reference architecture for the second institution

Highlights of their architecture include the following:

- Using on-premises virtual infrastructure for the Satellite location
- Using a proxy that they manage for traffic to and from the Satellite location
- Running Satellite-enabled Red Hat OpenShift on IBM Cloud for their workloads
- Integrating with various on-premises solutions (outside of the Satellite location) for logging, monitoring, image registry, and DevOps
- Using the following services on IBM Cloud:
  - IAM for identity and access management, federated with an external identity provider (Active Directory Federation Services)
  - Object Storage for Red Hat OpenShift on IBM Cloud control plane backups
  - Red Hat NTP to sync the time on all nodes

Next steps

- Satellite shared responsibility model
- Satellite architecture best practices

Satellite reference architecture shared responsibility model

In IBM Cloud, the responsibilities for deploying, operating, and securing products are shared between IBM and our customers. This IBM Cloud® shared responsibility model is one of the most important things to understand when interpreting the control requirements of the IBM Cloud Framework for Financial Services. Here we dive deeper into this shared responsibility model for the Satellite reference architecture.

Shared responsibilities for IBM Cloud products

All services that run in IBM Cloud as part of the Satellite reference architecture are considered managed products in the IBM Cloud shared responsibility model. Each of these products has their own service-specific shared responsibilities.

Due to the hybrid nature of Satellite, there is a unique set of Satellite-specific shared responsibilities for components that run in the on-premises Satellite location.

For all products, these responsibilities cut across the following five task areas, which intersect with the IBM Cloud Framework for Financial Services's best practices and requirements:

Types of tasks	Description
Incident and operations management	Includes tasks such as monitoring, event management, high availability, problem determination, recovery, and full state backup and recovery.
Change management	Includes tasks such as deployment, configuration, upgrades, patching, configuration changes, and deletion.
Identity and access management	Includes tasks such as authentication, authorization, access control policies, and approving, granting, and revoking access.
Security and regulation compliance	Includes tasks such as security controls implementation and compliance certification.

Disaster Recovery	Includes tasks such as providing dependencies on disaster recovery sites, provision disaster recovery environments, data and configuration backup, replicating data and configuration to the disaster recovery environment, and failover on disaster events.

Tasks areas in IBM Cloud shared responsibility model

Additional roles for shared responsibility matrix

The two-party view of responsibility doesn't adequately account for the broader set of roles (often represented by teams or even different companies) that need to collaborate to ensure that the overall solution provides a secure platform for workloads and data. In the following diagram and table, a broader set of roles is defined.

Roles for Satellite reference architecture RACI matrix

Role	Description
Workload consumer	Line-of-business organization that uses the workload or represents the internal or external users of the workload.
Workload provider	Client IT development and operations team responsible for developing, deploying, and managing the workload and user responsibility for the PaaS layer (for example, by providing updated operating system images or requesting worker node upgrades).
IBM Cloud on-premises PaaS provider	IBM Cloud development and operations teams responsible for the Satellite components deployed on-premises. Many supporting services, processes, and operations run in IBM Cloud.
On-premises IaaS / data center provider	Client IT operations and facilities teams responsible for the data center, networking, hardware, and virtualization that supports the on-premises PaaS and workload.
IBM Cloud IaaS / data center provider	IBM Cloud development and operations teams responsible for the direct Satellite management capabilities and supporting services that run in IBM Cloud.

Roles for deploying, managing, and operating all components of the Satellite reference architecture

Overview of shared responsibilities

Review the following table of who is responsible for particular cloud resources when using Satellite. In the table, "Shared" means that there is a shared responsibility between the workload provider and the IBM Cloud on-premises PaaS provider.

Resource	Incident and operations management	Change management	Identity and access management	Security and regulation compliance	Disaster Recovery
Client data	Workload provider	Workload provider	Workload provider	Workload provider	Workload provider
Application	Workload provider	Workload provider	Workload provider	Workload provider	Workload provider
Satellite Location	Shared	Shared	Shared	Shared	Shared
Satellite Host	Shared	Shared	Shared	Shared	Shared
Satellite Config	Shared	Shared	Shared	Shared	Shared
Satellite Link	Shared	Shared	Shared	Shared	Workload provider
Satellite Storage	Shared	Shared	Workload provider	Shared	Shared

Satellite-enabled services	Shared	Shared	Shared	Shared	Shared
Operating System	Workload provider	Shared	Workload provider	Shared	Workload provider
Virtual and bare metal servers	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider
Virtual storage	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider
Virtual network	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider
Hypervisor	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider
Physical servers and memory	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider
Physical storage	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider
Physical network and devices	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider
Facilities and data centers	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider	On-prem IaaS / data center provider

Overview of shared responsibilities.

Satellite location control plane and workload clusters

The following table goes a level deeper and shows the shared responsibilities for the control plane worker nodes and your workload clusters in the Satellite location. In the table, "Shared" means that there is a shared responsibility between the workload provider and the IBM Cloud on-premises PaaS provider.

Resource	Control plane worker nodes	Workload clusters
Client data	n/a	Workload provider
Application	Shared	Workload provider
Satellite Host	Shared	Shared
Satellite Config	Shared	n/a
Satellite Link	Shared [1]	n/a
Satellite Storage	Shared [2]	Shared [3]
Host networking	On-premises PaaS provider	Workload provider
Operating System	Shared [4]	Shared [5]
Virtual and bare metal servers	On-premises IaaS / data center provider	On-premises IaaS / data center provider
Virtual storage	On-premises IaaS / data center provider	On-premises IaaS / data center provider

Virtual network	On-premises IaaS / data center provider	On-premises IaaS / data center provider
Hypervisor	On-premises IaaS / data center provider	On-premises IaaS / data center provider
Physical servers and memory	On-premises IaaS / data center provider	On-premises IaaS / data center provider
Physical storage	On-premises IaaS / data center provider	On-premises IaaS / data center provider
Physical network and devices	On-premises IaaS / data center provider	On-premises IaaS / data center provider
Facilities and data centers	On-premises IaaS / data center provider	On-premises IaaS / data center provider

Overview of shared responsibilities in Satellite location.

Other on-premises components outside of the IBM Cloud Satellite location

The workload provider and on-premises IaaS / data center provider are solely responsible for edge plane, management plane, and other applications that you run in the on-premises environment outside of the Satellite location.

Next steps

- Satellite architecture best practices

---

1. Responsibility is shared, except that the workload provider is solely responsible for Disaster Recovery. ↩
2. Responsibility is shared, except that the workload provider is solely responsible for Identity and access management. ↩
3. Responsibility is shared, except that the workload provider is solely responsible for Identity and access management. ↩
4. Responsibility is shared, except that the workload provider is solely responsible for Change management and Security and regulation compliance. ↩
5. Responsibility is shared, except that the workload provider is solely responsible for Change management and Security and regulation compliance. ↩

Satellite reference architecture best practices checklist

IBM Cloud for VMware Regulated Workloads architecture

Another architectural option for the IBM Cloud for Financial Services is the IBM Cloud for VMware Regulated Workloads reference architecture. The architecture is an extension of the VMware vCenter Server® offering. Its design extends and enhances the basic vCenter Server architecture to deliver a secure, high-performance platform. You're able to run both classic virtualized workloads and containerized applications with the addition of Red Hat OpenShift on IBM Cloud.

Bring your own solutions for VPC reference architecture samples

Web application firewall and full-tunnel VPN

Deploying and configuring F5 BIG-IP

If you need to have a firewall within the VPC, it can be achieved using F5 BIG-IP. In this tutorial, you will learn how to provision an instance of BIG-IP. This is a prerequisite for using BIG-IP to set up a full tunnel client-to-site VPN or enable a WAF in a way that meets IBM Cloud Framework for Financial Services requirements.

⚠ **Important:** Guidance is provided here, but you are solely responsible for installing, configuring, and operating IBM third-party software in a way that satisfies IBM Cloud Framework for Financial Services requirements. In addition, IBM does not provide support for third-party software.

The architecture diagram shows a deployment of the VPC reference architecture with an instance of BIG-IP in the edge/transit VPC.

Single-region IBM Cloud for Financial Services reference architecture for VPC with BIG-IP

Tip: The architecture diagram shows a BIG-IP installation with four interfaces within 4 different VPC subnets: management, external, workload, and bastion. Depending on the services that you need and how you name your VPC subnets, you must update the values within the tutorial.

Objectives

The objective of this solution is to provision an instance of BIG-IP that is:

- Licensed with a bring-your-own license (BYOL).
- Configure management UI cipher suite.
- Replace the configuration utility's self-signed device certificate.

Before you begin

You need the following items to deploy and configure this reference architecture with an instance of BIG-IP:

- BIG-IP license
- Edge/transit VPC with a number of subnets that depend on the services that you plan to configure with the BIG-IP. All of these subnets must be in the same zone as your BIG-IP installation, and must be replicated for each zone BIG-IP is deployed to. Note of the ID's of your provisioned subnets because they are used later in this tutorial:
  - 3 subnets: If you are using only the BIG-IP full tunnel VPN service (external, management, bastion)
  - 3 subnets: If you are using only the BIG-IP WAF (external, management, workload)
  - 4 subnets: If you are using the BIG-IP full tunnel VPN and WAF service (external, management, workload, bastion)
- VPC SSH key
- Following best practices for access management and create an access group with minimum access policies for VPC Infrastructure Services:
  - Platform access: Editor
  - Service access: Writer and IP spoofing operator

Step 1: Provision BIG-IP

1. Install the F5 BIG-IP Virtual Edition for VPC offering from the IBM Cloud catalog. In the **Parameters without default values** section, pay special attention to the following parameters and the values that need to be set:
   - **external_subnet_id**: ID of the external subnet
   - **internal_subnet_id**: ID of the workload subnet
   - **tmos_admin_password**: Password for the admin account that must meet the following requirements:
     - Minimum length of 15 characters
     - Required Characters: Numeric = 1, Uppercase = 1, Lowercase = 1
2. In the **Parameters with default values** section, pay special attention to the following parameters and the values that should be set:
   - **cluster_subnet_id**: ID of the bastion subnet
   - **bigip_external_floating_ip**: `true`
   - **instance_profile**:
     - If you are using BIG-IP for either full tunnel VPN or web application firewall, use a minimum profile of `cx2-4x8`.

- If you are using BIG-IP for full tunnel VPN and web application firewall, use a minimum profile of `cx2-8x16`.
- **tmos_image_name**: `bigip-16-1`

3. To encrypt the boot disk of the F5 you must enter a value for the optional parameter encryption_key_crn . To get the value for that parameter:
   a. In the IBM Cloud console, go to your HPCS instance page
   b. On the left side of the page select 'KMS Keys'
   c. Find which key you use to encrypt disks. SLZ normally names this key "vsi-volume"
   d. Click on the three dot menu on the right side of the row containing your key
   e. Select View Key Details
   f. Copy the entire contents of the 'Cloud Resource Name' to your clipboard and use it to populate the value for `encryption_key_crn`
4. When you click **Install**, an IBM Cloud Schematics workspace is created and initialized.
5. Click **Generate Plan** in the top right.
6. After the plan is generated successfully, click **Apply Plan**.
7. Verify that the BIG-IP instance successfully deployed.

Note: Access the BIG-IP installation through the IP address only that is contained in the management subnet (https://<VSI-management-subnet-IP-address>). Use an access control list and security group to limit traffic of the BIG-IP management console to the IP address of the management subnet.

Step 2: Complete licensing for BIG-IP

1. Access BIG-IP configuration utility through your web browser, and login by using the *admin* user account.
2. Click the **Licensing** tab in the left menu, and then click **Activate** to complete licensing.
3. Provide the registration key, and click **Next**.
4. Click **Accept**. This restarts the BIG-IP modules and services and redirects to the login page.
5. Log in again, using the admin credentials.

Step 3: Configure management UI cipher suite

Important: You are responsible for meeting certain TLS requirements along with acceptable cipher suites.

Next, configure the management UI to allow the acceptable cipher suites of TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 and TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 only. You can do this configuration only with the command line interface.

1. Log on with SSH port 22 to the management IP address of the BIG-IP.
2. Issue the following commands:

```
tmsh modify /sys httpd ssl-ciphersuite 'ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-GCM-SHA256'
tmsh save /sys config
tmsh list /sys httpd ssl-ciphersuite
```

3. Logout of the command line interface.

Step 4: Replacing the configuration utility's self-signed device certificate

By default, the configuration utility is deployed with a self-signed certificate. You must replace the certificate according to the TLS requirements. For more information on replacing the BIG-IP self-signed certificate, see Replacing the Configuration utility's self-signed device certificate with a CA-signed device certificate.

Step 5: Accessing the F5 BIG-IP

All access, whether it is CLI with SSH or the configuration utility, should be through the bastion host. The following example shows how you can connect to the F5 BIG-IP configuration utility by using port forwarding with Teleport as a bastion host.

1. Create a Teleport role with the option **port_forwarding** set to **true**.
2. Log in to the bastion host with the tsh client.
3. Set up a port forwarding connection with the tsh client. The following example binds port 5000 on your local machine and allows access to the BIG-IP through the browser at *https://127.0.0.1:5000*.

```
tsh ssh -L 5000:<IP ADDRESS OF BIG-IP MANAGEMENT INTERFACE>:443 [user@]bastion_host
```

4. Access the BIG-IP configuration utility through your browser at `https://<bind_ip>:<listen_port>`. When you access the configuration utility, the browser receives a warning because you are accessing it through `bind_ip` that you set.

⚠ **Important:** Teleport does not provide session recordings through port forwarding. However, Teleport does provide audit events of the session that is associated with port forwarding, and the F5 BIG-IP does have an audit log of all administrative actions.

Step 6: Enabling GUI audit logging

By default, the audit logging for the GUI is disabled. Perform the following steps to enable it.

1. Access BIG-IP configuration utility through your web browser, and login by using the *admin* user account.
2. Under **System > Logs > Configuration**, click the **Options**.
3. Selec **Enable** for the GUI setting in the Audit Logging area.
4. Click **Update**.

Step 7: Setup BIG-IP routes

BIG-IP routes control the flow of traffic and determine which interface to use. By routing to the proper interface, you are able to use access controls lists (ACLs) and security groups to control traffic to and from the subnet and interface. Route traffic coming into the interface labeled *External* (for example, 10.5.40.5) and route it to your workload or application (for example, 10.40.0.0/18).

1. Go to **Network > Routes**.
2. Click **Add**.
3. In the **Properties** field complete the following information:
   - **Name**: Name that you want to use for the route
   - **Description**: Description you want to use for the route
   - **Destination**: IP prefix of your target workload or application (for example, 10.40.0.0)
   - **Netmask**: Netmask prefix of your target workload or application (for example, 255.255.192.0)
   - **Resource**: `Use Gateway`
   - **Gateway Address**: Gateway IP address of the interface that is labeled **Workload**. This is the second address in the CIDR range of the subnet labeled **Workload** (for example, 10.5.50.1).

Note: Routes are also needed for Bastion connectivity with full tunnel VPN that uses the interface that is labeled **Bastion** and for the F5 management UI that uses the interface that is labeled **MGMT**.

Step 8: Access control lists and security groups

You are responsible for isolating and controlling traffic from each VPC subnet and interface. By using access control lists and security groups, you can control all incoming and outgoing traffic.

Note: The example security groups and access and control list's based on the architecture diagram and the F5 interfaces (management, external, workload, bastion) in Zone 1 of the edge/transit VPC. The example assumes that the VPN is on port 4443 and your workload target port is 443. Include any other services that the F5 BIG-IP will need connectivity to. It is encouraged to update to meet your needs and architecture.

F5 management interface

Note: Restrict SSH port 22 and F5 configuration utility port 443 into the BIG-IP to the bastion host.

Security group

Protocol	Source type	Source	Port
TCP	CIDR Block	CIDR Block value from Bastion host subnet (for example, 10.5.70.0/24)	Ports 22-22
TCP	CIDR Block	CIDR Block value from Bastion host subnet (for example, 10.5.70.0/24)	Ports 443-443

F5 management interface: Security group inbound rules

Access control list

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	IP or CIDR	CIDR Block value from Bastion host subnet (for example, 10.5.70.0/24)	Any	IP address	IP address of F5 management interface (for example, 10.5.30.4)	22
Allow	TCP	IP or CIDR	CIDR Block value from Bastion host subnet (for example, 10.5.70.0/24)	Any	IP address	IP address of F5 management interface (for example, 10.5.30.4)	443
Deny	ALL	Any		Any	Any		

F5 management interface: access control list inbound rules

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	IP or CIDR	IP address of F5 management interface (for example, 10.5.30.4)	22	IP or CIDR	CIDR Block value from Bastion host subnet (for example, 10.5.70.0/24)	Any
Allow	TCP	IP or CIDR	IP address of F5 management interface (for example, 10.5.30.4)	443	IP or CIDR	CIDR Block value from Bastion host subnet (for example, 10.5.70.0/24)	Any
Deny	ALL	Any		Any	Any		

F5 management interface: access control list outbound rules

F5 external interface

Note: Restrict incoming traffic on port 443 to CIS global load balancers allowlisted IP addresses. The allowlisted IP addresses might periodically change and should be updated.

Security group

Protocol	Source type	Source	Port
TCP	CIS global load balancers allowlisted IP addresses		Ports 443-443
TCP	Any		Ports 4443-4443

F5 external interface: Security group inbound rules

Access control list

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	Any		Any	IP address	IP address of F5 external interface (for example, 10.5.40.5)	443
Allow	TCP	Any		Any	IP address	IP address of F5 external interface (for example, 10.5.40.5)	4443

Deny	ALL	Any		Any	Any

F5 external interface: access control list inbound rules

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	IP or CIDR	IP address of F5 external interface (for example, 10.5.40.5)	443	Any		Any
Allow	TCP	IP or CIDR	IP address of F5 external interface (for example, 10.5.40.5)	4443	Any		Any
Deny	ALL	Any		Any	Any		

F5 external interface: access control list outbound rules

F5 workload interface

Security group

Protocol	Destination type	Destination	Port
TCP	CIDR block	CIDR block of the subnet your workload or application is located (for example, 10.40.0.0/18)	Ports 443-443
TCP	CIDR block	CIDR block of the subnet your workload or application is located (for example, 10.50.0.0/18)	Ports 443-443
TCP	CIDR block	CIDR block of the subnet your workload or application is located (for example, 10.60.0.0/18)	Ports 443-443

F5 workload interface: Security group outbound rules

Access control list

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	IP or CIDR	CIDR block of the subnet your workload or application is located (for example, 10.40.0.0/18)	443	IP or CIDR	IP address of F5 workload interface (for example, 10.5.50.5)	Any
Allow	TCP	IP or CIDR	CIDR block of the subnet your workload or application is located (for example, 10.50.0.0/18)	443	IP or CIDR	IP address of F5 workload interface (for example, 10.5.50.5)	Any
Allow	TCP	IP or CIDR	CIDR block of the subnet your workload or application is located (for example, 10.60.0.0/18)	443	IP or CIDR	IP address of F5 workload interface (for example, 10.5.50.5)	Any
Deny	ALL	Any		Any	Any		

F5 workload interface: access control list inbound rules

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	IP or CIDR	IP address of F5 workload interface (for example, 10.5.50.5)	Any	IP or CIDR	CIDR block of the subnet your workload or application is located (for example, 10.40.0.0/18)	443

Allow	TCP	IP or CIDR	IP address of F5 workload interface (for example, 10.5.50.5)	Any	IP or CIDR	CIDR block of the subnet your workload or application is located (for example, 10.50.0.0/18)	443
Allow	TCP	IP or CIDR	IP address of F5 workload interface (for example, 10.5.50.5)	Any	IP or CIDR	CIDR block of the subnet your workload or application is located (for example, 10.60.0.0/18)	443
Deny	ALL	Any		Any	Any		

F5 workload interface: access control list outbound rules

F5 bastion interface

Security group

Protocol	Destination type	Destination	Port
TCP	CIDR block	CIDR Block value from Bastion subnet (for example, 10.5.70.0/24)	Ports 3023-3025
TCP	CIDR block	CIDR Block value from Bastion subnet (for example, 10.5.70.0/24)	Ports 3080-3080

F5 bastion interface: Security group outbound rules

Access control list

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	IP or CIDR	CIDR Block value from Bastion subnet (for example, 10.5.70.0/24)	3023-3025	IP or CIDR	IP address of F5 bastion interface (for example, 10.5.60.5)	Any
Allow	TCP	IP or CIDR	CIDR Block value from Bastion subnet (for example, 10.5.70.0/24)	3080-3080	IP or CIDR	IP address of F5 bastion interface (for example, 10.5.60.5)	Any
Deny	ALL	Any		Any	Any		

F5 Bastion interface: access control list inbound rules

Allow or Deny	Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
Allow	TCP	IP or CIDR	IP address of F5 bastion interface (for example, 10.5.60.5)	Any	IP or CIDR	CIDR Block value from Bastion subnet (for example, 10.5.70.0/24)	3023-3025
Allow	TCP	IP or CIDR	IP address of F5 bastion interface (for example, 10.5.60.5)	Any	IP or CIDR	CIDR Block value from Bastion subnet (for example, 10.5.70.0/24)	3080-3080
Deny	ALL	Any		Any	Any		

F5 bastion interface: access control list outbound rules

Next steps

- Set up a full tunnel client-to-site VPN with BIG-IP

- Enable a WAF with BIG-IP

Setting up a web application firewall with F5 BIG-IP

There are several architectural approaches to enabling consumer connectivity to the workload VPC. When using public internet access to the workload VPC, a web application firewall (WAF) is required. A WAF helps protect web applications by filtering and monitoring internet traffic. This tutorial guides you through one approach to enabling WAF by using IBM Cloud® Internet Services (CIS) and F5 BIG-IP Virtual Edition. Global load balancing along with a WAF forms the basis for you to meet IBM Cloud Framework for Financial Services requirements for boundary protection. Alternatively, IBM Cloud® Internet Services can independently fulfill the WAF requirements of the IBM Cloud Framework for Financial Services.

> ⚠ **Important:** Guidance is provided, but you are solely responsible for installing, configuring, and operating IBM third-party software in a way that satisfies IBM Cloud Framework for Financial Services requirements. In addition, IBM does not provide support for third-party software.

The following architecture diagram shows a deployment of the VPC reference architecture that uses BIG-IP for WAF.

Single-region IBM Cloud for Financial Services reference architecture for VPC with BIG-IP

There are two main components to this solution:

- CIS provides global load balancing and layer 3/4 protection against distributed denial-of-service (DDoS) attacks.
- BIG-IP that provides WAF protection and layer 7 protection against denial-of-service (DoS) attacks.

> **Note:** CIS is not Financial Services Validated. Because of this, TLS connections must not be terminated in CIS and should be configured only for pass-through connections. CIS global load balancers must be configured with the **proxy** configuration setting value of **off**. CIS Use range applications sto provide DDoS protection in front of global load balancers.

Network path

The following diagram shows the flow of network traffic from a client application passing through CIS and BIG-IP to server-side applications running in Red Hat OpenShift on IBM Cloud.

Network path from client application to Red Hat OpenShift on IBM Cloud application passing through CIS and BIG-IP.

The network flow goes through the following steps:

1. Client applications connect to a CIS Range application through a public hostname that is exposed by CIS. This connection is treated as a TCP connection (not an HTTPS connection) by CIS and passes through CIS without TLS termination. The Range application provides DDoS protection.
2. A CIS global load balancer routes traffic to the WAF provided by BIG-IP instances. The global load balancer has a public hostname, but it is not advertised for use.
3. The first TLS termination occurs in BIG-IP running in IBM Cloud. BIG-IP treats the connection as an HTTPS connection. The application's public certificate is used by BIG-IP for TLS termination.
4. Network traffic is reencrypted between BIG-IP and the application. The Red Hat OpenShift on IBM Cloud ingress controller is configured for pass through so that it does not terminate the TLS connection. The final TLS termination is done by the application.

The following table summarizes the fully qualified domain names (FQDNs) used by the three layers:

Host description	Example FQDN	DNS	How Protected
Range application hostname (user application hostname)	`www.application.example.com`	Managed by CIS, no explicit record	CIS Layer 3/4
GLB hostname	`<uuid>-example.com`	Managed by CIS, no explicit record, public network	Obfuscated, BIG-IP Virtual Server accepts only the Range Application Host Header
BIG-IP hostname	`f5-app.example.com`	A record maps to public external IP of BIG-IP, public network	Security Group to only allow traffic from CIS

Red Hat OpenShift on IBM Cloud application hostname (not externalized)	application-internal.example.com	CNAME record maps to private-router ALB, private network	Only accessible within network and protected by BIG-IP firewall. Security group restricts traffic to accept traffic only from BIG-IP.

Before you begin

The following items are needed to deploy and configure the reference architecture with CIS and BIG-IP to create a WAF:

- Provisioned and licensed BIG-IP
- Registered domain for CIS
- Certificate for your registered domain (can be a wildcard certificate)
- Intermediate certificate for your server-side application

Step 1: Provision CIS

1. Provision an instance of CIS that uses the **Enterprise Usage** plan.
2. Click **Create** to provision the service.

Step 2: Configure CIS TLS Mode

1. Register your domain name with CIS.
2. Go to **Security > TLS** set the TLS encryption mode to **End-to-end flexible**. **End-to-End CA signed** must not be used because it causes the TLS connection to be terminated in CIS.

Step 3: Configure CIS Global Load Balancer

When a global load balancer is created in CIS, it is automatically assigned the fully qualified domain name (FQDN) of `<global load balancer name>.<domain>`. This hostname is registered in public DNS. However, this hostname should not be used for the application because it lacks DDoS protection. Use he hostname that is assigned to the Range application in the next step for the application.

1. Go to **Reliability > Global load balancers**.
2. Click **Create**.
3. Set **Enable** to **On**.
4. Set **Proxy** to **Off**. You must do this to avoid terminating the TLS connection in CIS.
5. Enter a **Name** for the load balancer. This name is not used externally, but is referenced by the Range application.
6. Set values for **TTL** and **Traffic steering** that are appropriate for your application.
7. Click **Create**.
8. Configure an origin pool for the new load balancer that contains the BIG-IP endpoints.
9. Configure appropriate health checks.

You can use the following curl command to verify that the global load balancer is working correctly.

```
$ curl -v https://<global load balancer name>.<domain>/<path>
```

The verbose output from the curl command shows the signer of the SSL certificate that is terminating TLS. The signer should be the SSL certificate that is within the BIG-IP. If the signer of the certificate includes `Cloudflare`, then CIS is incorrectly terminating the TLS connection. Check that the global load balancer's **Proxy** is turned off and that the **TLS encryption mode** is set to **End-to-End CA Flexible**.

Step 4: Configure CIS Range Application

DDoS protection in CIS is enabled by using a Range application. The Range application becomes the application's front end and sets the FQDN for the service.

1. Go to **Security > Range**.
2. Click **Create**.
3. Select **TCP** as the Type. This enables pass-through without terminating TLS.
4. Set the **Name** to the hostname that you want to expose for the application. The application's FQDN becomes `<name>.<CIS domain>`.
5. Set the **Edge port** to the appropriate value. For HTTPS traffic, set the edge port to 443.
6. Select an appropriate value for **Edge IP connectivity**.
7. Set the **Origin** to **Load balancer**.
    a. Set **Load balancer** to the name of the global load balancer created in the previous step.

   b. Set the **Port** to the port that is exposed by the global load balancer.
8. You can enable **IP firewall** if you want to restrict the IP ranges that can connect to the service. If IP filtering is not needed, disable **IP firewall**.
9. **TLS Termination** must be set to **off** to prevent the TLS connection from being terminated in CIS.
10. Leave **TLS mode** greyed out.
11. Set **Proxy protocol** to **Off**.

You can use the following curl command to verify that the combination of the Range Application and Global Load Balancer is working correctly.

```
$ curl -v https://<Range application name>.<domain>/<path>
```

The verbose output from the curl command shows the signer of the SSL certificate where TLS is being terminated. The signer should be the signer of the application's SSL certificate. If the signer of the certificate includes "Cloudflare", then CIS is incorrectly terminating the TLS connection. Check that the Global Load Balancer's **Proxy** is configured to **Off** and that the **TLS encryption mode** is set to **End-to-End CA Flexible**.

When the edge application is working, it is important to prevent users from using the global load balancer directly. This is especially important to prevent a DDoS attack against the global load balancer. To prevent users from accessing the global load balancer directly, create a policy within BIG-IP that allows traffic to pass that specifies the FQDN of the Range application only.

Step 5: Enable license modules on the BIG-IP

1. Under **System > License**, click the **Module Allocation** tab.
2. If **Application Security (ASM)** is not enabled, enable it with the **Nominal** selection.
3. After BIG-IP restarts, log back into the BIG-IP console.

Step 6: Import client certificate and key to BIG-IP

Import client key

1. In the BIG-IP console, go to **System > Certificate Management > Traffic Certificate Management**.
2. Click **Import**.
3. Select **Key** from the **Import Type**.
4. Type in a name for your key under the option **Key Name**.
5. Either upload your key or paste your key under the option **Key Source**.
6. Select the proper security type.
7. Click **Import**.

Import client certificate

1. In the BIG-IP console, go to **System > Certificate Management > Traffic Certificate Management**.
2. Click **Import**.
3. Select **Certificate** from the **Import Type**.
4. In the **Certificate Name** field, enter a name for your certificate.
5. Either upload your certificate or paste your certificate under the option **Certificate Source**.
6. Click **Import**.

Step 7: Import server-side intermediate certificates to BIG-IP

For each server certificate that you have, complete the following actions:

1. In the console, go to **System > Certificate Management > Traffic Certificate Management**.
2. Click **Import**.
3. Select **Certificate** from the **Import Type**.
4. In the **Certificate Name** field, enter a name for your certificate.
5. Either upload your certificate or paste your certificate under the option **Certificate Source**.
6. Click **Import**.

Step 8: Configure ciphers

In this step, you create a cipher policy and group to limit the ciphers that are able to be used for a secure connection.

⚠ **Important:** For more information on TLS requirements and acceptable cipher suites, see Data encryption in transit.

Create a cipher policy

1. In the BIG-IP console, go to **Local Traffic > Ciphers > Rules** .
2. Click **Create**.
3. Set the following options:
   - **Name**: Name that you want to user for this policy
   - **Cipher Suites**: `ECDHE-ECDSA-CHACHA20-POLY1305-SHA256:ECDHE-RSA-CHACHA20-POLY1305-SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:TLSV1_3`
4. Click **Finished**.

Create a cipher group

1. In the BIG-IP console, go to **Local Traffic > Ciphers > Groups** .
2. Click **Create**.
3. Set the following options:
   - **Name**: Name that you want to user for this group
   - Under **Group Details**, move the rule that you created to the **Allow the following** section
4. Click **Finished**.

Step 9: Create SSL profiles in BIG-IP

In this step, you create a client and server SSL profile by using the client certificate or key and server certificate from the previous steps.

Create a client SSL profile

1. In the BIG-IP console, go to **Local Traffic > Profiles > SSL > Client** .
2. Click **Create**.
3. Set the following options:
   - **Name**: Name that you want to user for this profile
4. In the **Configuration** section, change **Basic** to **Advanced**, select **Custom** and set the following values:
   - **Certificate Key Chain**: Add both the key and certificate that you imported in the previous steps
   - **Ciphers**: select Cipher Group and the name of the group that you created in the previous step
   - **Options**: `Options List...`
   - **Options List**: The only enabled option should be `Don't insert empty fragments`
5. Click **Finished**.

Create a server SSL profile

For each server certificate that you imported, create a server SSL profile by completing the following steps:

1. In the BIG-IP console, navigate to **Local Traffic > Profiles > SSL -> Server** .
2. Click **Create**.
3. In the **General Properties** section, specify:
   - **Name**: Name that you want to use for this profile
4. In the **Configuration** section, change **Basic** to **Advanced**, click the **Custom** checkbox and set the following values:
   - **Secure Renegotiation**: **Require Strict**
   - **Server Name**: FQDN of your custom application
   - **Ciphers**: select Cipher Group and the name of the group you created in the step above
   - **Options**: `Options List...`
   - **Options List**: Only enabled option should be `Don't insert empty fragments`
5. In the **Server Authentication** section, set the following options:
   - **Server Certificate**: **require...**
   - **Trusted Certificate Authorities**: Set to your intermediate server certificate
6. Click **Finished**.

Step 10: Setting Up the HPCS integration with F5

Transport Layer Security (TLS) secures the communication between clients and servers, safeguarding against possible hacking attempts and man-in-the-middle attacks. TLS offloading involves employing a Hardware Security Module (HSM) to handle TLS encryption and decryption tasks instead of relying on the web server. This approach notably diminishes the risk of key compromise since TLS encryption keys are stored within FIPS 140-2 Level 4 validated HSMs. F5 Big IP offers configuration options through through which it can be integrated with HPCS for TLS offloading. To configure the HPCS to manage encryption keys and perform cryptographic operations in Big IP F5, see F5 docs on Setting Up the Network HSM.

Step 11: Create a BIG-IP virtual server

In this step, you create a BIG-IP virtual server for traffic distribution with a backend pool.

Create iRules

An iRule is a feature within the BIG-IP local traffic management (LTM) system that you can use to manage your network traffic. See Getting started with iRules: Basic Concepts for more information about iRules.

1. If you have a backend that serves multiple HTTPS sites by using the TLS Server Name Indication feature, for more information see Configure a virtual server to serve multiple HTTPS sites using the TLS Server Name Indication feature.
2. Host header rewrite is used if your public facing domain is different than your backend. Use the following example to rewrite the host header based on the FQDN of the incoming request:

```
when HTTP_REQUEST {
    set hostname [getfield [HTTP::host] ":" 1]
}
when HTTP_REQUEST_SEND {
    switch -glob [string tolower $hostname] {
    "mysiteA.public.domain" {
      clientside {
        HTTP::header replace host "mysiteA.backend.domain"
      }
    }
    "mysiteB.public.domain" {
      clientside {
        HTTP::header replace host "mysiteB.backend.domain"
      }
    }
    default {
      reject
    }
  }
}
```

Create the backend pool

1. In the BIG-IP console, go to **Local Traffic > Pools > Pool List** .
2. Click **Create**.
3. In the **Configuration** section, set the following fields:
    - **Name**: Name that you want to use for the pool
    - **Health Monitor**: Select an appropriate protocol to monitor the health of your members
4. In the **Resources** section, set the following fields:
    - **Load Balancing Method**: Select an appropriate method that you would like to use for balancing traffic
    - **New Members**: For each IP or FQDN, input the members of your backend server

Create policies

If you have multiple backend pools to which you need to route traffic, you must create a policy to route traffic to the proper backend pool. For more information, see Introducing Local Traffic Policies.

1. Go to **Local Traffic > Policies > Policy List** .
2. Click **Create**.
3. Set the following under **General Properties**:
    - **Policy Name**: Name that you want to use for the policy
    - **Strategy**: Set the strategy that you would like to use for rule matching

    - **Type**: **Traffic Policy**
4. Click **Create**.
5. Under **Rules**, click **Create**.
6. Under **Name**, enter the name that you want to specify for this rule.
7. Click the **+** symbol and enter the conditions that must be matched.
8. Click the **+** symbol and enter what should be done when the condition is met. If you want to forward to a backend pool, select **Forward Traffic**.
9. If you need to create multiple rules, click **Create** again and repeat the previous steps.
10. After all rules are created, click **Save Draft**.
11. Under **Local Traffic > Policies > Policy List**, select your policy and click **Publish**.

Create the virtual server

1. In the console, go to **Local Traffic > Virtual Servers > Virtual Server List** .
2. Click **Create**
3. In the **General Properties** section, select or input the following fields:
    - **Name**: Name that you want to use for the virtual server
    - **Type**: **Standard**
    - **Destination Address/Mask**: Internal IP of the External interface
    - **Service Port**: `443`
    - **State**: **Enabled**
4. Under the **Configuration** section, select or input the following fields:
    - **Protocol**: **TCP**
    - **Protocol Profile (Client)**: **tcp**
    - **Protocol Profile (Server)**: **(Use Client Profile)**
    - **HTTP Profile (Client)**: **http**
    - **HTTP Profile (Server)**: **(Use Client Profile)**
    - **SSL Profile (Client)**: Name of the client SSL profile you created previously
    - **SSL Profile (Server)**: If you have one server profile, select the name that you created in a previous step, otherwise see the iRules sections
    - **Source Address Translation**: **Auto Map**
5. Click **Finished**.
6. Click the name of your newly created virtual server from the **Virtual Server** List.
7. Click the **Resources** tab.
8. Under the **Load Balancing** section select the following options:
    - **Default Pool**: Name of backend pool that you created in a previous step
9. Under the **iRules** section:
    - Click **Manage**.
    - Move the iRules that you want to enable to the **Enabled** section.
10. Under the **Policies** section:
    - Click **Manage**.
    - Move the policy that you want to enable to the **Enabled** section.

Step 12: Configure WAF on the virtual server

Enable license

1. Under **System > License**, click the **Module Allocation** tab.
2. If **Application Security (ASM)** is not enabled, enable it with the **Nominal** selection.
3. BIG-IP restarts.
4. When the restart completes, log back into BIG-IP.

Set up WAF

1. Go to **Security > Guided Configuration**.
2. Click **Web Application Protocol** and then **Web Application Comprehensive Protection**.
3. Click **Next**.

4. Under **Security Layers*** do the following:
    a. Enter a configuration name.
    b. Enable any security layers that you would like. The security layers of *Security Policy* and *Behavioral Analysis DoS* should at least be enabled.
    c. Click **Save** and **Next**.
5. Under **Web Application Security Policy Properties**:
    a. Set the following:
    - **Select Enforcement Mode**: **Blocking**
    - **Select type of policy to protect application** : **Generic**
    - **Application Language**: Select the language encoding for your web application
    a. Click **Save** and **Next**.
6. Under **Bot Defense Properties**
    a. Set the following:
    - **Select Enforcement Mode**: **Blocking**
    - **Profile Template**: **Balanced**
    a. Click **Save** and **Next**.
7. Under **DoS Profile Properties**
    a. Set the following under **Operation Mode**:
    - **Operation Mode**: **Blocking**
    - Verify that **Bad actors behavior detection** and **Request signature detection** are selected
    a. Set the following under **Mitigation Mode**:
    - Mitigation Mode: **Standard**
    a. Click **Save** and **Next**.
8. Under **Virtual Server Properties**
    a. Check **Assign Policy to Virtual Server(s)**.
    b. Under **Virtual Server** select **Use Existing**.
    c. Under **Assign Virtual Servers**, move the virtual server that you created previously over from **Available** to **Selected**.
    d. Click **Save** and **Next**.
9. Click **Deploy**.
10. Verify that the configuration is under * *Deployed* status.

Setting up full tunnel client-to-site VPN that uses F5 BIG-IP

One of the approaches for enabling connectivity to the management VPC is with a client to site full tunnel VPN. Application providers who want to use this option must deploy their own solution. This tutorial shows one approach that uses F5 BIG-IP Virtual Edition that can be used to meet IBM Cloud Framework for Financial Services requirements. Alternatively, Client VPN for VPC can be used which also meets the requirements of the IBM Cloud Framework for Financial Services.

The following diagram shows a deployment of the VPC reference architecture with a BIG-IP full tunnel VPN deployment.

Single-region IBM Cloud for Financial Services reference architecture for VPC with BIG-IP full tunnel VPN host

Objectives

The objective of this tutorial is to create a full VPN tunnel host that uses BIG-IP that provides the following functions:

- LDAP authentication
- Access only to resources on the bastion subnet, including the bastion host

⚠ **Important:** BIG-IP currently cannot be integrated with App ID for authentication. Work is ongoing with the BIG-IP team to identify a solution.

Before you begin

The following items are needed to deploy and configure the reference architecture with a full tunnel VPN provided by BIG-IP:

- Provisioned and licensed BIG-IP
- Configured LDAP

Step 1: Enable license modules on the BIG-IP

1. Under **System > License**, click the **Module Allocation** tab.
2. If **Access Policy (APM)** is not enabled, enable it with the **Nominal** selection.
3. After BIG-IP restarts, log back into the BIG-IP console.

Step 2: Set up an AAA Server in BIG-IP for VPN users

For this setup, we are using an LDAP database for VPN users. BIG-IP provides several other connectors for authentication.

1. Log in to the BIG-IP configuration utility that uses the admin user account.
2. Go to **Access > Authentication > LDAP**.
3. Click **Create**.
4. Enter a name, for example, `vpn_users` and click **OK**.
5. For the **Server Connection** setting, select **Use Pool** even if you have only one LDAP server.
6. In the **Server Pool Name** field, type a name for the AAA server pool.
7. Populate the **Server Addresses** field by typing the IP address of a pool member and clicking **Add**.
8. For the **Mode** setting, select **LDAPS**.
9. In the **Service Port** field, retain the default port number for LDAPS, 636, or type the port number for the SSL service on the server.
10. In the **Admin DN** field, type the distinguished name (DN) of the user with administrator rights. Type the value in this format:
    `CN=administrator,CN=users,DC=sales,DC=mycompany,DC=com`
11. In the **Admin Password** field, type the administrative password for the server.
12. In the **Verify Admin Password** field, retype the administrative password for the server.
13. From the **SSL Profile (Server)** list, select an SSL server profile. You can select the default profile, **serverssl**, if you do not need a custom SSL profile.
14. Skip the **Timeout** field because it does not apply when you select **Use Pool**.
15. Click **Finished**.
16. The new LDAPS server displays on the LDAP server list.

For more information about LDAP setup as a AAA Server in BIG-IP, see the LDAP and LDAPS Authentication.

Step 3: Set up a client SSL profile

> **Note:** If you have set up and configured the web application firewall, you can skip to the next step.

1. Import client certificate and key to BIG-IP.
2. Configure cipher suites.
3. Create a client SSL profile.

Step 4: Create a BIG-IP VPN

1. GO to the **Wizards > Device Wizards** page.
2. Select **Network Access Setup Wizard for Remote Access** and click **Next**.
3. Provide a policy name and caption, for example, **vpn_policy** and click **Next**.
4. On the Authentication page, select **Use Existing, LDAP** for server type and select the LDAP server that you previously created. Provide a user DN, for example, `uid=%{session.logon.last.username},ou=Users,dc=example,dc=com`, and click **Next**.
5. Configure the **Lease Pool** by providing an IP range from the previously allocated subnet for VPN in the VPC subnets. For example, `10.5.10.0/24`. Click **Next**.
6. Keep the compression setting set to **No compression**. For the Traffic options, ensure that **Force all traffic through tunnel** is selected. Click **Next**.
7. Keep the DNS servers set to what is populated by default and click **Next**.
8. Provide the virtual server IP to be the private IP address of the external interface, for example, `10.5.40.4`. Click **Next**.
9. Review the wizard inputs and click **Next**.
10. Confirm the summary and click **Finished**.

Step 5: Configure the VPN virtual server

1. After the creation of VPN, navigate to **Local Traffic > Virtual Servers** and click the name of the virtual server that is associated with the VPN.
2. If you are using the web application service, under **Service Port**, change the port number to another number other than 443 (for example, 4443).
3. Under **SSL Profile (Client)**, move the client SSL profile that you created to the section labeled **Selected** and move any other ones out from under it to

the section labeled **Available**.

4. Click **Update**.

Step 6: Test the BIG-IP VPN

1. Access the BIG-IP WebTop by navigating to `https://&#60;floating-ip-address-on-external-interface-of-f5&#62;:&#60;port` that you defined for the VPN virtual server> in your browser.
2. Follow the prompts to allow BIG-IP endpoint inspection. You might be required to install the BIG-IP endpoint inspector application/plugin, if it is not already installed.
3. Enter the credentials for an LDAP user and click **Submit**.
4. Follow the prompts to allow BIG-IP VPN link. You might be required to install the BIG-IP VPN application/plugin, if it's not already installed.
5. A BIG-IP VPN connection details window displays.
6. You are now connected to a full tunnel VPN.
    - You have access to resources on the bastion subnet such as the bastion host.
    - You do not have access to resources on the workload subnet. The workload resources must be accessed by using the bastion.

Bastion host

Setting up a bastion host that uses Teleport

This tutorial shows you one way that can be used to meet the IBM Cloud Framework for Financial Services requirements that are related to bastion host. There are various ways to implement a compliant bastion solution, but we show you how to configure a bastion host in your VPC by using Teleport Enterprise Edition, along with Object Storage and App ID for enhanced security. You will learn how to set up a Teleport-based solution that meets the previously described IBM Cloud Framework for Financial Services requirements.

> ⚠ **Important:** We provide guidance, but you are solely responsible for installing, configuring, and operating IBM third-party software in a way that satisfies IBM Cloud Framework for Financial Services requirements. In addition, IBM does not provide support for third-party software. So, if you choose to use Teleport and encounter issues with the Teleport software that require support, you should contact Teleport.

To implement your bastion host solution, you must complete the following high-level steps:

1. Provision an instance of Object Storage and configure a bucket for storing session recordings.
2. Provision an instance of App ID and configure a SAML-based identity provider.
3. Provision a virtual server instance, install Teleport, and complete the Teleport configuration for unified access to your infrastructure.

Before you begin

Before you deploy and configure a bastion host into a virtual server instance running on the VPC reference architecture, make sure you have completed the following prerequisites:

- Acquire a Teleport Enterprise Edition license
- Configure a VPC like *VPC 0 (Edge Network)* in the architecture diagram

Single-region IBM Cloud for Financial Services reference architecture for VPC with bastion host

Limitations

The following limitations currently apply for the bastion host solution that is described:

- Windows is not supported.
- RDP is not supported.

Step 1: Configure Object Storage

Complete the following steps to configure Object Storage.

1. Provision a IBM Cloud® Object Storage instance with the `Standard Plan`.
2. Create a custom bucket for storing session recordings with the following configuration:
    - Storage class set to *Standard*
    - Resiliency set to *Regional*
    - Location set to the *region of your VPC*

- Retention set to *12 months*

3. Generate a new set of service credentials with the bucket permissions of *Object Write* access and the advanced option of *Include HMAC Credentials* enabled.

> ⚠ **Important:** Teleport session recordings can contain sensitive information. You should have policies to ensure that operators do not enter secrets on the command line interface. In addition, you need to employ least privilege and make sure that access to these recordings is restricted only to those employees who need the recordings to do their jobs (such as your security incident response team).

Step 2: Configure App ID

1. Provision an IBM Cloud® App ID instance with the `Graduated tier Plan`.
2. If you have your own SAML-based identity provider, complete the following steps:
   a. Configure your SAML-based identity provider within App ID.
   b. Set up MFA by using a physical hardware-based security key and your identity provider.
   c. Set up the following password policies:
   - Lock user accounts after three (3) consecutive failed logon attempts within 15 minutes.
   - Lock user accounts for 30 minutes after more than three unsuccessful logon attempts. When the lockout period ends, the user can reset their password. Internal privileged accounts must remain locked until they are released by an administrator.
   - Set sessions to time out after 15 minutes of inactivity.
3. Generate a service credential with the role of `Reader`.
4. From your App ID service dashboard, click **Manage Authentication > Authentication Settings** and add the web direct URLs of `https://<virtual server instance FQDN>:3080/v1/webapi/oidc/callback`.

Step 3: Provision virtual server instances for Teleport Enterprise

1. Provision a virtual server instance within each zone of the VPC cluster.
   - Profile: cx2-4x8 with 4 vCPUs, 8 GB RAM, 8 Gbps.
   - Linux-based operating system (CentOS, Debian, RHEL, Ubuntu).
2. If your identity provider is hosted publicly and not accessible directly from the Teleport virtual server instance in the VPC, you must provision a public gateway.
3. Register each virtual server instance that you plan to run Teleport into a Domain Name Server.
4. Generate an SSL certificate for each of the provisioned virtual server instances or use a wildcard certificate.

Step 4: Configure a security group for your virtual server instances

Configure the Security group for your virtual server instances.

Protocol	Source type	Source	Port
TCP	IP address or CIDR Block	IP address or CIDR Block value	Ports 22-22
TCP	IP address or CIDR Block	IP address or CIDR Block value	Ports 3023-3025
TCP	IP address or CIDR Block	IP address or CIDR Block value	Ports 3080-3080

Bastion host: Security group inbound rules

Protocol	Destination type	Destination	Value
TCP	IP address or CIDR Block	IP address or CIDR Block value	Ports Any

Bastion host: Security group outbound rules

> **Note:** Add any additional security group rules that are needed for connectivity to your resources.

Step 5: Configure an access control list

Configure the access control list for the subnets that the bastion virtual server instances use to `Allow` inbound and outbound traffic. Make sure that there is a rule for each bastion virtual server instance if the ACL is used for multiple subnets.

Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
TCP	IP address or CIDR Block	IP address or CIDR Block value	Any	IP address	IP address of bastion server	22
TCP	IP address or CIDR Block	IP address or CIDR Block value	Any	IP address	IP address of bastion server	3023-3025
TCP	IP address or CIDR Block	IP address or CIDR Block value	Any	IP address	IP address of bastion server	3080

Bastion host: Access control list inbound rules

Protocol	Source type	Source Value	Source Port	Destination type	Destination Value	Destination Port
TCP	IP address	IP address of bastion server	22	IP address or CIDR Block	IP address or CIDR Block value	Any
TCP	IP address	IP address of bastion server	3023-3025	IP address or CIDR Block	IP address or CIDR Block value	Any
TCP	IP address	IP address of bastion server	3080	IP address or CIDR Block	IP address or CIDR Block value	Any

Bastion host: Access control list outbound rules

Note: Add any additional access control list rules that are needed for connectivity to your resources.

Step 6: Configure virtual server instances for Teleport Enterprise

With your virtual server instances provisioned and set-up with a security group and ACL, you can continue configuring the instance and installing Teleport. Then, complete the Teleport configuration for unified access to your infrastructure.

1. Connect to your virtual server instance by using SSH.
2. Download Teleport Enterprise version 7.1.0 and extract its contents. (Optionally, you can follow the instructions to download the TAR file for the latest Teleport Enterprise version). Copy the following packages to the `/usr/local/bin` directory:
   - `cp <path of extracted contents>/teleport /usr/local/bin`
   - `cp <path of extracted contents>/tctl /usr/local/bin`
   - `cp <path of extracted contents>/tsh /usr/local/bin`
3. Create the directory `/var/lib/teleport` on the file system of your virtual server instance. This directory might exist depending on your installation method.
4. Copy your license file from Teleport Enterprise into the file `/var/lib/teleport/license.pem`.
5. Copy your SSL certificate for the virtual server instance into `/var/lib/teleport/`. You should have a file for the certificate and key. If you have an intermediate certificate, make sure it is after your certificate.
6. Create the file `teleport.yaml` in the directory `/etc` and copy the sample content from the following example into the file. Ensure the appropriate changes where noted with `<variables>`.

```
#teleport.yaml
teleport:
  nodename: <fqdn of node>
  data_dir: /var/lib/teleport
  log:
    output: stderr
    severity: DEBUG
  storage:
    audit_sessions_uri: "s3://<Bucket>?endpoint=<COS Endpoint>&region=ibm"
```

```
auth_service:
  enabled: "yes"
  listen_addr: 0.0.0.0:3025
  authentication:
    type: oidc
    local_auth: false
  license_file: /var/lib/teleport/license.pem
  message_of_the_day: "<banner message to be displayed to a user that must be acknowledged before logging into the bastion>"

ssh_service:
  enabled: "yes"
  commands:
  - name: hostname
    command: [hostname]
    period: 1m0s
  - name: arch
    command: [uname, -p]
    period: 1h0m0s

proxy_service:
  enabled: "yes"
  listen_addr: 0.0.0.0:3023
  web_listen_addr: 0.0.0.0:3080
  tunnel_listen_addr: 0.0.0.0:3024
  https_cert_file: /var/lib/teleport/<SSL Certificate PEM File>
  https_key_file: /var/lib/teleport/<SSL Certificate Key PEM File>
```

7. Create the file `/etc/systemd/system/teleport.service` with the following example content. Make sure to enter your HMAC credentials that you generated previously in this procedure. For more information, see Systemd Unit File. Ensure the appropriate changes where noted with `<variables>`.

```
[Unit]
Description=Teleport Service
After=network.target

[Service]
Type=simple
Restart=on-failure
Environment=AWS_ACCESS_KEY_ID="<HMAC access_key_id>"
Environment=AWS_SECRET_ACCESS_KEY="<HMAC secret_access_key>"
ExecStart=/usr/local/bin/teleport start --config=/etc/teleport.yaml --pid-file=/run/teleport.pid
ExecReload=/bin/kill -HUP $MAINPID
PIDFile=/run/teleport.pid

[Install]
WantedBy=multi-user.target
```

8. Issue the following systemctl to load the teleport service.

```
$ sudo systemctl daemon-reload
sudo systemctl start teleport
sudo systemctl enable teleport
```

9. Some operating systems have open firewalls but if your operating system limits traffic, you must add firewall rules to allow TCP traffic for ports 3023, 3024, 3025, and 3080. CentOS is an example of an operating system that blocks traffic by default.

10. Start the Teleport process `systemctl start teleport`.

11. Create a Teleport Role. For more information, see Create Teleport Roles. For more information on settings for a Teleport role, see Teleport Access Control Reference.

    a. Create the file `role.yaml` within the directory `/var/lib/teleport`. The following sample role provides full access to the system. You can also create roles within the Teleport web console under *Teams -> Roles*.

```
#example role
kind: "role"
version: "v3"
metadata:
  name: "teleport-admin"
spec:
```

```
options:
  max_connections: 3
  cert_format: standard
  client_idle_timeout: 15m
  disconnect_expired_cert: no
  enhanced_recording:
  - command
  - network
  forward_agent: true
  max_session_ttl: 1h
  port_forwarding: false
allow:
  logins: [root]
  node_labels:
    "*": "*"
  rules:
  - resources: ["*"]
    verbs: ["*"]
```

12. Create the OIDC connector.

    a. Create the file `oidc.yaml` within the directory `/var/lib/teleport` by using the following example content. Ensure the appropriate changes where noted with `<variables>`. For more information, see OIDC Authentication.

```
#oidc connector
kind: oidc
version: v2
metadata:
  name: appid
spec:
  redirect_url: "https://<virtual server instance DNS FQDN>:3080/v1/webapi/oidc/callback"
  client_id: "<Client ID from AppID Service Credentials>"
  display: AppID
  client_secret: "<secret from AppID Service Credentials>"
  issuer_url: "<oauthServerUrl from AppID Service Credentials>"
  scope: ["openid", "email"]
  claims_to_roles:
  - {claim: "email", value: "<Email Address>", roles: ["teleport-admin"]}
```

    Example claim names can be `email`, `family_name`, `given_name`, or `name`. The value is what that claims value will be set to.

13. Using the tctl to apply the yamls:

    - `tctl create /var/lib/teleport/role.yaml`
    - `tctl create /var/lib/teleport/oidc.yaml`

14. Set up forwarding of Teleport logs and system logs. Teleport logs are located in the directory `/var/lib/teleport` and system logs in `/var/logs`.

> ⚠ **Important:** Logs must be forwarded to an operational logging solution. For more information, see operational logging

Step 7: Log in to the bastion host

You can log in to the bastion host through the web console or tsh client as described in the following sections.

Log in through the the web console

1. Access the web console on port 3080.

   `https://<fqdn of node>:3080`

2. Start a terminal session under *Servers*. There should be a single server with a *connect* button. Click *connect* and select the user that you would like to log in with.

Log in through the tsh client

1. Install the Teleport client tool tsh.
2. Log in using tsh.

```
tsh login --proxy=<fqdn of telport server>:3080
```

3. Run shell or execute a command on a remote SSH node by using the tsh ssh command

```
tsh ssh <[user@]host>
```

Step 8: Install tools

Now that the bastion host is set up and configured, you can install the tools that are needed to interact with your infrastructure, such as:

1. IBM Cloud CLI and associated plug-ins.
2. OpenShift Origin CLI.

Note: For information on accessing your Red Hat OpenShift on IBM Cloud cluster via the tools above, see accessing Red Hat OpenShift clusters.

Step 9: Remove SSH port from security group and access control list

Now that Teleport is installed and setup, remove SSH port 22 from the allowed list of ports within the configured security group and access control list that is assigned to the virtual server and subnet.

Related controls in IBM Cloud Framework for Financial Services

See the related controls for bastion host.

General deployment and configuration information

Set up environment for deployment and configuration

Automation is an important part of any cloud solution, and it's even more so in regulated industries. You want to ensure that your deployment, operations, and management procedures are secure and repeatable. Manual activities can be error prone and lead to the introduction of vulnerabilities. So, as you work with the reference architectures for the IBM Cloud for Financial Services, you will want to take advantage of the tools that IBM Cloud provides to make automation possible.

Tip: If using Virtual Private Cloud (VPC) reference architecture, is highly recommended that you use the VPC landing zone deployable architectures which provide a preconfigured set of infrastructure as code (IaC) assets to help you get started with your deployments.

The remainder of this topic describes how to set up the following tools for working with your IBM Cloud resources.

- IBM Cloud Command Line Interface (CLI)
- Application programming interfaces (APIs)
- Terraform for IBM Cloud

Important: Use private routes to IBM Cloud service endpoints to enhance control and security over your data when you use the services that make up the reference architectures. Private routes are not accessible or reachable over the internet. By using them, you can protect your data from threats from the public network and logically extend your private network. For more information, see Secure access to services by using service endpoints and Securing your connection when you use the IBM Cloud CLI .

IBM Cloud Command Line Interface

The CLI offers a powerful set of commands to work with your resources. Most of the services that are part of the reference architectures have specific plug-ins that you can use to extend the base CLI experience.

The following table provides links for the CLI extensions for each service in the reference architectures that has one.

Category	VPC reference architecture	Satellite reference architecture	Optional for both
Core	• VPC infrastructure services [1]	• Satellite	
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry	• Red Hat OpenShift on IBM Cloud [2] • Container Registry	
Networking	• VPC infrastructure services • Direct Link • Transit Gateway		
Storage	• Block Storage for VPC • Object Storage	• Object Storage	
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services	• App ID [3]

Logging and monitoring	• Activity Tracker Event Routing [4] • Compliance Manager [5] • Flow Logs for VPC [6] • IBM Cloud Logs	• Activity Tracker Event Routing • Compliance Manager • IBM Cloud Logs	
Integration			• Event Streams

CLI information for services in reference architectures

Application programming interfaces

IBM Cloud offers a rich set of APIs for working with your resources. The following table provides links for the APIs that can be used for each service in the reference architectures.

Category	VPC reference architecture	Satellite reference architecture	Optional for both
Core	• VPC infrastructure services [7]	• Satellite	
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry	• Red Hat OpenShift on IBM Cloud [8] • Container Registry	
Networking	• VPC infrastructure services • Direct Link • Transit Gateway		
Storage	• Block Storage for VPC • Object Storage	• Object Storage	
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services	• App ID
Logging and monitoring	• Activity Tracker Event Routing • Compliance Manager • Flow Logs for VPC • IBM Cloud Logs	• Activity Tracker Event Routing • Compliance Manager • IBM Cloud Logs	
Integration			• Event Streams

API information for services in reference architectures

Terraform for IBM Cloud

Terraform on IBM Cloud enables predictable and consistent provisioning of IBM Cloud resources so that you can rapidly build IaC to deploy complex cloud environments. The following table contains links about the Terraform capabilities for each service in the reference architectures.

Category	VPC reference architecture	Satellite reference architecture	Optional for both

Core	• VPC infrastructure services [9]	• Satellite
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry	• Red Hat OpenShift on IBM Cloud [10] • Container Registry
Networking	• VPC infrastructure services • Direct Link • Transit Gateway	
Storage	• Block Storage for VPC • Object Storage	• Object Storage
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services • App ID
Logging and monitoring	• Activity Tracker Event Routing [11] • Compliance Manager • Flow Logs for VPC • IBM Cloud Logs	• Activity Tracker Event Routing • Compliance Manager • IBM Cloud Logs
Integration		• Event Streams

Terraform information for services in reference architectures

Next steps

- Deploy infrastructure as code for reference architectures

1. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
2. Satellite-enabled service which runs in your Satellite location. ↩
3. No specific CLI information is available for App ID. ↩
4. No specific CLI information is available for Activity Tracker Event Routing. ↩
5. No specific CLI information is available for Compliance Manager. ↩
6. No specific CLI information is available for Flow Logs. ↩
7. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
8. Satellite-enabled service which runs in your Satellite location. ↩
9. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
10. Satellite-enabled service which runs in your Satellite location. ↩
11. No specific Terraform information is available for Activity Tracker Event Routing. ↩

Use deployable architectures to deploy your resources

Infrastructure as Code (IaC) automates the provisioning of infrastructure and enables your organization to develop, deploy, and scale cloud applications

with greater speed, less risk, and reduced cost. With that in mind, IBM Cloud offers a set of *deployable architectures* that automate the deployment of different variations of the VPC reference architecture.

See Account setup in IBM Cloud

Intra Organization Multitenancy Architecture Reference

As explained in Organizing IBM Cloud accounts and resources, each deployment of one of the reference architectures is single tenant, which means the deployment is intended for a single customer. However, there are some cases where you might need different business units in the same organization to be managed by the same operator.

See Best practices for multitenant management in IBM Cloud to get started.

The guide referenced above provides different deployment models focused mostly on multitenancy for container-based solutions. A combination of these models can be used in a solution deployment, for example, critical and non-critical environments can have different models.

Also consider additional separation for different environment types, such as using different clusters and infrastructure for different environment types when using namespace isolation.

Next steps

- Account setup in IBM Cloud

Accounts, identity management, and access control

IBM Cloud account setup

An IBM Cloud account is needed to provision and manage IBM Cloud services that make up the reference architectures of the IBM Cloud for Financial Services. Along with the high-level steps to follow, we describe some of the best practices for account setup that will help you satisfy the requirements of the IBM Cloud Framework for Financial Services. In addition, the most relevant control requirements are provided.

1. Create an IBM Cloud account. For more information, see Create your account.

   > **Tip:** It is highly recommended that you use a functional ID that is owned by your company rather than an employee's personal ID. A functional ID is a company-owned email address (such as `ibm-cloud-admin@domain.com`) used to represent a functional user. This allows for uninterrupted administrative access by the account owner as employees leave the company or are reassigned to other projects.

   The following table shows the controls that are most related to this step.

Family	Control
Access Control (AC)	AC-2 Account Management

Related controls in IBM Cloud Framework for Financial Services for account creation [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management

Related controls in IBM Cloud Framework for Financial Services for account creation [FSv1.1]

2. Set up the Activity Tracker Event Routing service as described in Audit logging for IBM Cloud events. This enables IBM Cloud platform events to be recorded for auditing purposes. Setting this up early in the process is important so that all platform events that occur during the rest of these steps are available in the audit logs.

   The following table shows the controls that are most related to this step.

Family	Control
Access Control (AC)	AC-2 Account Management AC-2 (1) Account Management | Automated System Account Management AC-2 (4) Account Management | Automated Audit Actions AC-2 (7) Account Management | Privileged User Accounts
Audit and Accountability (AU)	AU-3 Content of Audit Records AU-4 Audit Log Storage Capacity AU-5 Response to Audit Logging Process Failures AU-6 Audit Record Review, Analysis, and Reporting AU-6 (1) Audit Record Review, Analysis, and Reporting | Automated Process Integration AU-7 Audit Record Reduction and Report Generation AU-10 Non-repudiation AU-11 Audit Record Retention

Related controls in IBM Cloud Framework for Financial Services for audit logging [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management AC-2 (1) Account Management | Automated System Account Management AC-2 (4) Account Management | Automated Audit Actions AC-2 (7) Account Management | Privileged User Accounts

Audit and Accountability (AU)	AU-3 Content of Audit Records AU-4 Audit Log Storage Capacity AU-5 Response to Audit Processing Failures AU-6 Audit Record Review, Analysis. and Reporting AU-6 (1) Audit Record Review, Analysis. and Reporting | Automated Process Integration AU-7 Audit Record Reduction and Report Generation AU-10 Non-repudiation AU-11 Audit Record Retention

Related controls in IBM Cloud Framework for Financial Services for audit logging [FSv1.1]

3. Upgrade your account to either Pay-As-You-Go or Subscription. For more information, see Upgrading your account.

> **Tip:** It is highly recommended to upgrade to a Subscription account so that you can set up an enterprise. Enterprises offer significant advantages in your ability to scale your environment over time as described in Enterprise account architecture.

4. Enable multi-factor authentication (MFA) for all users in your account. Choose MFA devices that align with the defined requirements. This can include hardware tokens (U2F), like FIDO2-compliant security keys, smart cards, or software tokens (TOTP). FIDO U2F standard offers the highest level of security.

   The following table shows the controls that are most related to this step.

Family	Control
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (organizational Users) | Multi-factor Authentication to Privileged Accounts IA-2 (6) Access to Accounts — Separate Device

Related controls in IBM Cloud Framework for Financial Services for multi-factor authentication [FSv2.0]

Family	Control
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (Organizational Users) | Multi-factor Authentication To Privileged Accounts IA-2 (11) Identification and Authentication (Organizational Users) | Remote Access - Separate Device

Related controls in IBM Cloud Framework for Financial Services for multi-factor authentication [FSv1.1]

5. Restrict IP addresses from which a user can access the IBM Cloud account. For more information, see Allowing specific IP addresses for an account for more information.

   The following table shows the controls that are most related to this step.

Family	Control
Access Control (AC)	AC-4 Information Flow Enforcement
System and Communications Protection (SC)	SC-7 Boundary Protection SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception

Related controls in IBM Cloud Framework for Financial Services for restricting IP addresses [FSv2.0]

Family	Control
Access Control (AC)	AC-4 Information Flow Enforcement
System and Communications Protection (SC)	SC-7 Boundary Protection SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception

Related controls in IBM Cloud Framework for Financial Services for restricting IP addresses [FSv1.1]

6. While optional, it is recommended that you enable authentication from an external identity provider (IdP) to securely authenticate external users to your IBM Cloud account. This provides a way for your employees to use your company's single sign-on (SSO) solution.

7. Enable the IBM Cloud for Financial Services Validated setting in your account. With this setting, you can filter the catalog for services that are designated as Financial Services Validated and indicates that your account stores regulated financial services information. If you enable Financial

Services Validated, your account still has access to the full public catalog. For more information, see Enabling your account to use Financial Services Validated products.

The following table shows the controls that are most related to this step.

Family	Control
Access Control (AC)	AC-20 Use of External Systems
System and Services Acquisition (SA)	SA-4 Acquisition Process SA-9 External System Services
Enterprise System and Services Acquisition (ESA)	ESA-5 Subcontractor Risk Management
Security Assessment and Authorization (CA)	CA-3 Information Exchange

Related controls in IBM Cloud Framework for Financial Services for using only Financial Services Validated services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
System and Services Acquisition (SA)	SA-4 Acquisitions Process SA-9 External Information System Services
Enterprise System and Services Acquisition (ESA)	ESA-5 Subcontractor Risk Management
Security Assessment and Authorization (CA)	CA-3 System Interconnections

Related controls in IBM Cloud Framework for Financial Services for using only Financial Services Validated services [FSv1.1]

8. Set the session inactivity timeout to 15 minutes. For more information, see Setting the sign-out due to inactivity duration.

   The following table shows the controls that are most related to this step.

Family	Control
Access Control (AC)	AC-11 Device Lock

Related controls in IBM Cloud Framework for Financial Services for using only session inactivity timeout [FSv2.0]

Family	Control
Access Control (AC)	AC-11 Session Lock

Related controls in IBM Cloud Framework for Financial Services for using only session inactivity timeout [FSv1.1]

9. Update company profile details.

   The following table shows the controls that are most related to this step.

Family	Control
Configuration Management (CM)	CM-8 (4) System Component Inventory | Accountability Information

Related IBM Cloud Framework for Financial Services controls for updating company profile details [FSv2.0]

Family	Control
Configuration Management (CM)	CM-8 (4) Information System Component Inventory | Accountability Information

Related IBM Cloud Framework for Financial Services controls for updating company profile details [FSv1.1]

10. Set email preferences for notifications. You can receive email notifications about IBM Cloud platform-related items, such as announcements, critical events, security notices, billing and usage, and ordering.

The following table shows the controls that are most related to this step.

Family	Control
System and Information Integrity (SI)	SI-2 Flaw Remediation SI-5 Security Alerts, Advisories, and Directives

Related IBM Cloud Framework for Financial Services controls for configuring notifications [FSv2.0]

Family	Control
System and Information Integrity (SI)	SI-2 Flaw Remediation SI-5 Security Alerts & Advisories

Related IBM Cloud Framework for Financial Services controls for configuring notifications [FSv1.1]

11. Choose a support plan. For more information, see Basic, Advanced, and Premium Support plans.

Next steps

- Organizing your IBM Cloud accounts and resources

Organizing IBM Cloud accounts and resources

After you complete your IBM Cloud account setup, consider the best practices for organizing your accounts and resources. The choices you make have a significant impact as you set up access management in later steps, as well as how well your organization can scale over time.

Each deployment of one of the reference architectures is single tenant, meaning the deployment is intended to be used by a single consumer. We recommend creating a new account for each deployment. If you have multiple consumers, then you need multiple accounts. Multiple individual accounts might quickly become unwieldy, so we recommend that you use an enterprise account.

Within an account, all resources managed by IBM Cloud Identity and Access Management (IAM) are placed into resource groups for access control and billing purposes.

For more information, see:

- Best practices for organizing resources and assigning access
- Managing resource groups

Organization for a single deployment within an account

VPC reference architecture

For the VPC reference architecture, it is recommended that you use two resource groups for each deployment:

- Management - Holds all of the resources that are needed by your management VPC
- Workload - Holds all of the resources that are needed by your workload VPC

This creates a clean separation between the layers. The following diagram shows such a layout.

Account with two resource groups for management and workload VPCs

Note: While not shown in the diagram, it is also permissable to take isolation a step further and place your management VPC in one account and your workload VPC in another account.

Satellite reference architecture

For the Satellite reference architecture, you can put all of your IBM Cloud resources that are related to a single Satellite deployment into one resource group.

Account with one resource group for Satellite-related resources

Organization for multiple deployments

It is required that you have separate test and production deployments of the reference architecture, and that those deployments are managed by separate accounts. This means you need at least two accounts, each with a dedicated resource group.

VPC reference architecture

The following diagram shows a depiction of this organization for the VPC reference architecture.

Multiple accounts for VPC reference architecture with two resource groups

Satellite reference architecture

The following diagram shows a depiction of this organization for the Satellite reference architecture.

Multiple accounts for Satellite reference architecture

Scaling the number of deployments within an account

If for some reason you do not want to use an enterprise account and you have a limited number of consumers, then you might also consider a variation where you maintain separate resource groups for each deployment within an account.

VPC reference architecture

For the VPC reference architecture, this variation is shown in the following diagram. In this case, we still have a test and production account, but we create two new resource groups for each deployment.

Using two accounts with multiple resource groups for VPC reference architecture

While this reduces the number of accounts you need, it can also become problematic because each account is subject to quotas and service limits for VPC.

Satellite reference architecture

For the Satellite reference architecture, we still have a test and production account. We then create one new resource group for each deployment as shown in the following diagram.

Using two accounts with multiple resource groups for Satellite reference architecture

Managing scalability with an enterprise

Rationale for using an enterprise

Unless you have a limited number of deployments and accounts, it is highly recommended that you use an enterprise. Large enterprises that allow an account structure, cross account networking, resource deployment, and billing to develop organically run the risk of encountering governance, scaling, security, and accounting issues. See Enterprise account architecture for recommendations on how to address these concerns across accounts so that a robust, compliant, and scalable solution can be achieved.

Organizing your enterprise

Within an enterprise, you can create multiple accounts and account groups. With this structure, you can easily manage many accounts and many deployments.

VPC reference architecture

The following diagram shows a single enterprise with one account group that contains separate accounts for each deployment of the VPC reference architecture.

Enterprise account organization with VPC reference architecture

When using enterprise accounts, it is recommended that you use one enterprise account for production and another enterprise account for development. See Planning your enterprise account structure for more information.

Satellite reference architecture

The following diagram shows an enterprise with one account group that contains separate accounts for each deployment of the Satellite reference architecture.

	Enterprise account organization

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-5 Separation of Duties AC-6 Least Privilege

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-5 Separation of Duties AC-6 Least Privilege

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Access management in IBM Cloud

Access management in IBM Cloud

After deciding how to organize your accounts and resources, you will need to properly manage access for deployments of the reference architectures that will help achieve separation of duties and least privilege.

Managing access with access groups

An access group is a grouping of user and service IDs to which the same IAM access can be granted. You can assign a single policy to the group instead of assigning the same access multiple times per individual user or service ID. A logical way to assign access is by creating one access group per wanted level of access.

⚠ **Important:** Always use access groups for managing access. IAM policies should never be attached directly to users.

For more information:

- IAM access for definitions of platform management roles and service access roles.
- Giving access to resources in resource groups
- Best practices for assigning access for more details about working with IAM.

You must determine what set of access groups works well for your situation. Whatever structure you choose, repeat it for every account (whether a stand-alone account or part of an enterprise).

Example access group structure

For illustration, the following table shows one possible setup that provides for a reasonable separation of duties.

Access Group	Description
cloud-organization-admins	Responsible for organizing the structure of the resources used by the organization.
cloud-network-admins	Responsible for creating networks, VPCs, load balancers, subnets, firewall rules, and network devices.
cloud-security-admins	Responsible for establishing and managing security policies for the entire organization, including access management and organization constraint policies.

cloud-billing-admins	Responsible for setting up billing accounts and monitoring their usage.
cloud-devops	DevOps practitioners create or manage end-to-end pipelines that support continuous integration and delivery, monitoring, and system provisioning.
cloud-developers	Developers are responsible for designing, coding, and testing applications.

Access groups

After the access groups are created, you can assign roles to them.

Tip: Never assign roles directly to users. Assign roles to access groups only.

For more information, see:

- IAM roles and actions for a list of applicable roles for every IBM Cloud service
- Access management in IBM Cloud

Invite users to your account

After you have access groups in place, then you can start inviting users to your account and assigning them to the appropriate access group depending on their job responsibilities. For more information, see Inviting users to an account.

Access groups with dynamic rules

If you're using an external identity provider, dynamic rules allow you to automatically add federated users to an access group based on SAML assertions. When a user logs in with a federated ID, the data that is provided by your identity provider dynamically maps the user to an access group based on the rules that you set. This can dramatically ease the administration of your account.

For more information see:

1. Creating dynamic rules for access groups

Considerations for enterprises

If you're using an enterprise, these resources provide additional information:

- Setting up an enterprise

Required permissions for services in reference architecture

The following table provides references to additional information for managing access with IAM for each service in the reference architectures.

Category	VPC reference architecture	Satellite reference architecture	Optional for both
Core	• VPC infrastructure services [1]	• Satellite	
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry	• Red Hat OpenShift on IBM Cloud [2] • Container Registry	
Networking	• VPC infrastructure services • Direct Link • Transit Gateway		

Storage	• Block Storage for VPC • Object Storage	• Object Storage	
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services	• App ID
Logging and monitoring	• Activity Tracker Event Routing • Compliance Manager • Flow Logs for VPC	• Activity Tracker Event Routing Compliance Manager	
Integration			• Event Streams

Managing access for IBM Cloud services in the reference architectures

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-3 Access Enforcement AC-5 Separation of Duties AC-6 Least Privilege

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-3 Access Enforcement AC-5 Separation of Duties AC-6 Least Privilege

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Handling and securing secrets

1. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
2. Satellite-enabled service which runs in your Satellite location. ↩

Handling and securing secrets

A secret is any piece of data that is sensitive within the context of an application or service. Secrets must be securely protected through their entire lifecycle.

Secrets include all of the following but are not limited to:

- Passwords of any type (database logins, OS accounts, functional IDs, and so on)
- API keys
- Long-lived authentication tokens (OAuth2, GitHub, IAM, and so on)
- SSH keys
- Encryption keys
- Other private keys (PKI/TLS certificates, HMAC keys, signing keys, and so on)

Application providers should ensure:

- Secrets are generated and stored in the environment (for example, dev, test, and production) where your service is deployed.
- Secrets never leave their environments (for example, dev, test, and production) and should be secured by using access control measures. Service design should minimize the number of machines and people with access to secrets by using both authorization and network restrictions based on the principle of least privilege.
- Secrets are rotated in according with the requirements of the IBM Cloud Framework for Financial Services with minimal or no downtime.
- Secrets are never stored in source code, configuration files, or documentation.

The following table lists the different solutions that you can use to protect your application secrets.

Scenario	What to use
You need to create, lease, and manage API keys, credentials, database configurations, and other secrets for your services and applications.	Use IBM Cloud Secrets Manager.
You need to generate, renew, and manage TLS/SSL certificates for your deployments.	Use IBM Cloud Secrets Manager.
You need to create and manage encryption keys.	Use Hyper Protect Crypto Services to manage encryption keys in a single-tenant service with dedicated hardware.
You need secrets in your Red Hat OpenShift on IBM Cloud environment for microservices to connect to system resources.	Use Kubernetes secrets encrypted using Hyper Protect Crypto Services as your Key Management Service (KMS) provider.

Secrets management and data protection scenarios

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-2 Account Management
Identification and Authentication (IA)	IA-2 Identification and Authentication (organizational Users) IA-5 Authenticator Management

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management
Identification and Authentication (IA)	IA-2 User Identification and Authentication IA-5 Authenticator Management

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Consumer accounts for application provider workloads

Consumer accounts for application provider workloads

You must have a system for authentication and authorization of the consumer organization's users when they connect to your application workloads through a web app or API. As the best practice for enabling a zero trust environment says, you need to provide proper role-based access control (RBAC) for these users.

You can use App ID to secure your apps, back-end resources, and APIs by using standards-based authentication. App ID makes it easy to add an authentication step to your applications with a few lines of code. You can add email or username, social, or enterprise sign in to your apps with APIs, SDKs, prebuilt UIs, or your own branded UIs.

You can choose between several identity providers. For more information, see Managing authentication. The two most likely options that you would use for

IBM Cloud for Financial Services are:

- Security Assertion Markup Language (SAML) - You can create a single sign-on experience for your users by integrating with the consumers identity provider.
- Cloud Directory - You can maintain your own user registry in the cloud. When a user signs up for your app, they are added to your directory of users. This option gives your users more freedom to manage their own account within your app.

App ID integrations with IBM Cloud services

You can use App ID with other IBM Cloud offerings. For example, if you're using Red Hat OpenShift on IBM Cloud, you can configure Ingress in your cluster to secure your apps at the cluster level. For more details, see Setting up Ingress and App ID authentication Ingress annotation to get started.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you have met the requirements.

Family	Control
Access Control (AC)	AC-2 Account Management AC-5 Separation of Duties AC-6 Least Privilege
Identification and Authentication (IA)	IA-5 Authenticator Management IA-5 (1) Authenticator Management | Password-based Authentication

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management AC-5 Separation of Duties AC-6 Least Privilege
Identification and Authentication (IA)	IA-5 Authenticator Management IA-5 (1) Authenticator Management | Password-Based Authentication

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Tagging IBM Cloud resources and managing access through tags

Tagging IBM Cloud resources and managing access with tags

All IBM Cloud resources in your account should be tagged with security attributes that you define. Apply user tags to organize your resources and easily find them later, to help you with identifying specific team usage or cost allocation, and to control access to your resources without requiring updates to your IAM policies.

For more information, see:

- Working with tags (through the Global Search and Tagging platform service)
- Controlling access to resources by using tags

Example tagging scheme

The IBM Cloud Framework for Financial Services doesn't require a specific set of tags that must be used, but they should be meaningful for your application. For illustration, you might use a scheme such as the one outlined.

In this example, it is suggested that you use one the following impact levels from Federal Information Processing Standards Publication 199:

- `fips199-high`: The loss of confidentiality, integrity, or availability could be expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, or individuals.
- `fips199-moderate`: The loss of confidentiality, integrity, or availability could be expected to have a serious adverse effect on organizational operations, organizational assets, or individuals.

- `fips199-low`: The loss of confidentiality, integrity, or availability could be expected to have a limited adverse effect on organizational operations, organizational assets, or individuals.

In addition, for services that store data, it is suggested that you use one or more of these tags corresponding to different kinds of data:

- `audit`: Audit logs or archives
- `public`: Publicly accessible, such as Cloud Object Storage buckets that hold public documentation or samples
- `consumer-owned-data`: Data owned by the consumer
- `consumer-metadata`: Metadata owned by the consumer
- `regulated-data`: Data that is regulated

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-16 Security and Privacy Attributes
System and Communications Protection (SC)	SC-16 Transmission of Security and Privacy Attributes

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-16 Security Attributes
System and Communications Protection (SC)	SC-16 Transmission of Security Attributes

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

If using the VPC reference architecture, then see:

- VPC networking overview

If using the Satellite reference architecture, then see:

- Satellite networking overview

Networking

Networking for VPC reference architecture

VPC networking overview

Properly setting up and managing your networking components is critical to building a solution that meets the best practices and requirements of the IBM Cloud Framework for Financial Services. This is the starting point for learning how to perform those tasks for the VPC reference architecture.

Next steps

Walk through the topics for how to build out your networking solution:

- Creating and connecting the management and workload VPCs
- Connecting application provider to the management VPC
- Running operator actions through bastion host
- Consumer connectivity to workload VPC
- Connectivity to IBM Cloud services with private endpoints
- Accessing the public internet
- FQDN based egress proxy in IBM Cloud VPC
- Connecting on prem environment with service providers on IBM Cloud with Direct Link

Creating and connecting the management and workload VPCs

After completing the work for account setup and management, you can now create the management and workload VPCs from the VPC reference architecture and connect them using Transit Gateway.

1. Create two VPCs in a multizone region, one for the management VPC and another for the workload VPC. See Create a VPC for more details. For now, you should not follow the instructions in any of the other sections in that reference.

   Tip: Do *not* create default address prefixes when you create the VPCs. You can specify `manual` for the `--address-prefix-management` argument in the ibmcloud is vpc-create command, such as in `ibmcloud is vpc-create my-vpc --address-prefix-management manual`.

2. Review Designing an addressing plan for a VPC to get guidance how to plan for addressing within your VPC. VPC uses Classless Inter-Domain Routing (CIDR) notation for specifying addresses.

3. Create the address prefix for each zone in both VPCs based on the plan that you developed in the previous step. See ibmcloud is vpc-address-prefix-create for details.

4. Create the subnets for the three zones by using the CIDRs. For more information, see Create a subnet.

   Tip: You need to specify zones. For `us-south`, the three zones are `us-south-1`, `us-south-2`, `us-south-3`. For `us-east`, the three zones are `us-east-1`, `us-east-2`, `us-east-3`. See multizone regions for more information, including the zone identifiers for other multizone regions.

5. Create security groups to define inbound and outbound traffic that's allowed for virtual server instances. For more information, see Using security groups and Overview of network security options.

6. Configure ACLs for your subnets that do not contain virtual servers or Red Hat OpenShift on IBM Cloud clusters (for example, subnets that contain a VPN Gateway or VPEs). For more information, see Set up network ACLs.

   Even though we recommend security groups where possible, there are subnets without virtual servers or Red Hat OpenShift on IBM Cloud clusters that require ACLs. In our example, we're referring to the subnets that contain virtual private endpoints (VPEs) only or VPN for VPC.

   For more information, see:

   - Configuring ACLs and security groups for use with VPN

7. Order Transit Gateway.

   For additional consideration, see the following resources:

   - Interconnect two or more VPCs in the same MZR
   - Planning for Transit Gateway
   - Team based privacy that uses IAM, VPC, Transit Gateway, and DNS

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-20 Use of External Systems
Security Assessment and Authorization (CA)	CA-3 Information Exchange
System and Communications Protection (SC)	SC-2 Separation of System and User Functionality SC-3 Security Function Isolation SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
Security Assessment and Authorization (CA)	CA-3 System Interconnections
System and Communications Protection (SC)	SC-2 Application Partitioning SC-3 Security Function Isolation SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Application provider connectivity to management VPC

Connecting application provider to the management VPC

The management VPC should be accessed only by you, the application provider. It's important that the connection be secure to avoid bad actors gaining access and conducting malicious operations. There are two options to enable this connectivity from your on-premises enterprise network: Direct Link and VPN for VPC. Alternatively, if you want to support connectivity without going through your enterprise network, you can deploy your own full tunnel client-to-site VPN solution. After a connection is established, operators can complete actions through a bastion host in the management VPC.

Note: Operators who are connecting to the on-premises enterprise network from offsite (such as their home) should connect to the enterprise network only by using a full tunnel client-to-site VPN solution. After connected to the enterprise network through a full tunnel, they can access the management VPC to perform their duties.

Direct Link

Direct Link is the most secure way to enable connectivity from the application provider's on-premises environment to the management VPC. The speed and reliability of Direct Link extends your organization’s data center network and offers more consistent, higher-throughput connectivity, keeping traffic within the IBM Cloud network. When using Direct Link, a private Application Load Balancer for VPC (ALB) is used to distribute traffic among multiple server instances within the same region of your VPC.

The following diagram shows the Direct Link connection pattern.

Application provider on-prem to management VPC using Direct Link

For more information, see:

- Interconnecting your VPC that uses IBM Cloud offerings
- Getting started with Direct Link
- Ordering IBM Cloud Direct Link
- Integrating with Virtual Private Endpoint for VPC

VPN for VPC

An alternative connectivity pattern is to use the VPN for VPC service to securely connect from your private network to the management VPC. VPN for VPC can be used as a static, route-based VPN or a policy-based VPN to set up an IPsec site-to-site tunnel between your VPC and your on-premises private network, or another VPC.

The following diagram shows the VPN for VPC connection pattern.

Application provider on-prem to management VPC using VPN for VPC

For more information, see:

- About VPN gateways
- Creating a VPN gateway
- Use a VPC/VPN gateway for secure and private on-premises access to cloud resources

Full tunnel client-to-site VPN

The third option for connectivity for your operators is to use a full tunnel client-to-site VPN, so they do not have to be on your on-premises network. However, IBM does not provide a Financial Services Validated full tunnel client-to-site VPN solution. So, if you want to use this option, you need to deploy your own. See Setting up full tunnel VPN with FS BIG-IP for one example of how to do this.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-4 (21) Information Flow Enforcement | Physical or Logical Separation of Information Flows AC-17 Remote Access AC-20 Use of External Systems
Audit and Accountability (AU)	AU-10 Non-repudiation
Security Assessment and Authorization (CA)	CA-3 Information Exchange
System and Communications Protection (SC)	SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-10 Network Disconnect SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-4 (21) Information Flow Enforcement | Physical or Logical Separation of Information Flows AC-17 Remote Access AC-20 Use of External Systems

Audit and Accountability (AU)	AU-10 Non-repudiation
Security Assessment and Authorization (CA)	CA-3 Information Exchange
System and Communications Protection (SC)	SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-10 Network Disconnect SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Performing operator actions through a bastion host

Running operator actions through a bastion host

All interactive operator actions must be run through a bastion host in the management VPC. A bastion host is a server that can be accessed through SSH, Windows Remote Desktop Protocol (RDP), or `kubectl`, but only through a Direct Link or VPN for VPC connection. After set up, the bastion host allows a secure connection to virtual server instances or Red Hat OpenShift on IBM Cloud clusters within the management VPC and the workload VPC. Administrative tasks on the individual servers are completed by using SSH, RDP, or `kubectl`, proxied through the bastion. Access to the servers and regular internet access from the servers, for example, for software installation, is allowed only with a special maintenance security group that is attached to those servers. In addition, session auditing must be enabled to record all privileged user actions.

After connecting to the management VPC through Direct Link or VPN for VPC, a complete bastion solution should include the following details:

- Full session recording of actions executed via the bastion to perform session audits. This includes SSH, and Kubernetes exec (`kubectl exec -it`) sessions.
  - Linux-based session recordings are captured as a raw dump of stdout and stderr streams, including TAB characters (bash escape sequences).
  - Session recordings should be placed in to a storage service maintained within your environment for archiving. The storage should be configured as “immutable object storage.” Retention policies are should be applied to the storage, so that data is stored in a WORM (Write-Once-Read-Many), non-erasable and non-rewritable manner. The policy is set and enforced for a 12-month (minimum) retention period.
- Authenticating users with MFA using a physical hardware-based security key that generates a six-digit numerical code. A smart card or hardware token designed and operated to FIPS 140-2 level 2 or above or equivalent (e.g., ANSI X9.24, ISO 13491-1:2007) is recommended.
- Locking user accounts after three (3) consecutive failed logon attempts within 15 minutes.
- Locking user accounts for 30 minutes when there have been more than three unsuccessful logon attempts. After the lockout period ends, the user will be able to reset their password. Internal privileged accounts must remain locked until released by an administrator.
- Session timeout after 15 minutes of inactivity.
- Providing a system use notification banner from either the bastion or the target system. The warning banner is displayed before the system grants access to the user and the usage conditions must be approved before proceeding. The banner will provide privacy and security notices consistent with applicable customer policies, regulations, standards, and guidance. The warning banner will state that:
  - Users are accessing a financial services information system.
  - Information system usage may be monitored, recorded, and subject to audit.
  - Unauthorized use of the information system is prohibited and subject to criminal and civil penalties.
  - Use of the information system indicates consent to monitoring and recording.

Set up a bastion host

You need to install and manage your own bastion solution within your management VPC. There are various ways a bastion solution can be implemented. For one example that uses Teleport Enterprise Edition, see Setting up a bastion host for secure connectivity .

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control

Audit and Accountability (AU)	AU-14 Session Audit
Access Control (AC)	AC-6 (9) Least Privilege | Log Use of Privileged Functions AC-17 Remote Access
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (organizational Users) | Multi-factor Authentication to Privileged Accounts

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Audit and Accountability (AU)	AU-14 Session Audit
Access Control (AC)	AC-6 (9) Least Privilege | Auditing Use of Privileged Functions AC-17 Remote Access
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (Organizational Users) | Network Access to Privileged Accounts

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Consumer connectivity to workload VPC

Consumer connectivity to workload VPC

Previously, we saw how administrative access to the workload VPC can be accomplished from the bastion host in the management VPC. Now, we look at how consumers can connect to the workload VPC to access your service offering.

Consumer in same organization as application provider

If the consumer is in the same organization that you are (such as the same financial institution), then the connection options are much the same as they are when you connect to the management VPC. That is, the consumer can connect to the workload VPC with either Direct Link or VPN for VPC. This is shown in the diagram below.

Connecting to workload VPC from on-prem with consumer in same organization as application provider

Direct Link

Direct Link is the most secure way to enable connectivity from the consumer's on-premises environment to the workload VPC. The speed and reliability of Direct Link extends your organization’s data center network and offers more consistent, higher-throughput connectivity, keeping traffic within the IBM Cloud network. When using Direct Link, a private Application Load Balancer for VPC (ALB) is used to distribute traffic among multiple server instances within the same region of your VPC.

The following diagram shows the Direct Link connection pattern.

Consumer on-prem to workload VPC using Direct Link

For more information, see:

- Interconnecting your VPC by using IBM Cloud offerings
- Getting started with Direct Link

VPN for VPC

An alternative connectivity pattern requires using the VPN for VPC service to securely connect from your private network to the management VPC. VPN for VPC can be used as a static, route-based VPN or a policy-based VPN to set up an IPsec site-to-site tunnel between your VPC and your on-premises private network, or another VPC.

The following diagram shows the VPN for VPC connection pattern.

Consumer on-prem to workload VPC using VPN for VPC

For more information, see:

- About VPN gateways
- Creating a VPN gateway

Consumer in different organization than application provider

Connecting from private intranet

Workloads may need to be accessed from private intranets using VPN or Direct Link connectivity setups. While Client-to-Site VPN and Site-to-Site VPN can handle NATing, using Direct Link to provide access to the workload requires special considerations to avoid IP space conflicts. Both parties—the consumer and the provider—need to coordinate the appropriate IP ranges for the VPC. The recommended approach is to connect the consumer's Direct Link to a Transit VPC, which hosts a firewall device that provides NATing capabilities, thereby enabling access to the workload.

Consumer on-prem to workload VPC using Direct Link when consumer and provider are different organizations

Connecting from public internet

There are many valid cases where you might want to allow consumers to access your service through the public internet. The base architecture can be adapted to securely enable this type of access as shown in the following diagram which introduces a new edge VPC. The request from the consumer gets routed through Cloud Internet Service's global load balancer, through a public load balancer in the edge VPC, and then to the private load balancer within the workload VPC. This is shown in the following diagram.

Detailed VPC reference architecture with edge VPC

Global load balancer

One option for global load balancing outside of the edge VPC is IBM Cloud® Internet Services (CIS), powered with Cloudflare. CIS provides a fast, highly performant, reliable, and secure internet service for customers running their business on IBM Cloud.

For more information, see the following resources:

- Getting started with CIS
- Configuring a global load balancer
- Protecting TCP traffic (Range)

Edge VPC with web application firewall

The edge VPC is used to enhance boundary protection for both the management VPC and the workload VPC. For public internet access to the workload VPC, a WAF in the (CIS)](/docs/cis?topic=cis-getting-started) is used to protect web applications by filtering and monitoring internet web traffic. A WAF can prevent attacks exploiting a web application's known vulnerabilities.

For management VPC connectivity, your operators can connect to the environment from your on-premises network (with Direct Link or VPN for VPC). In practice, all three zones in the edge VPC would be the same, but for illustrative purposes, each of the first and second zone in the edge VPC box depicts one of the two scenarios for operator connectivity:

- Zone 1 - Connectivity with Direct Link, so a VPN for VPC is not needed.
- Zone 2 - Connectivity from the application provider is with VPN for VPC, so a Direct Link is not needed.

Finally, the bastion can be put either in the edge VPC or the management VPC. If you place it in the edge VPC, then the management VPC becomes optional if you are not deploying any other management tools.

Application load balancer in workload VPC

Use IBM Cloud® Application Load Balancer for VPC (ALB) to distribute traffic among multiple server instances within the same region of your VPC. For more information, see the following resources:

- About IBM Cloud Application Load Balancer for VPC
- Creating an Application Load Balancer for VPC

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
System and Communications Protection (SC)	SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
System and Communications Protection (SC)	SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Connectivity to IBM Cloud services with private endpoints

Connectivity to IBM Cloud services with private endpoints

IBM Cloud services should be used only over private routes. Private routes are not accessible or reachable over the internet. By using the IBM Cloud private endpoints feature, you can protect your data from threats from the public network and logically extend your private network.

When inside a VPC, this private access can be accomplished by using a virtual private endpoint (VPE) to map a VPC IP address to the IBM Cloud service. VPEs are virtual IP interfaces that are bound to an endpoint gateway created on a per service, or service instance, basis (depending on the service operation model). The endpoint gateway is a virtualized function that scales horizontally, is redundant and highly available, and spans all availability zones of your VPC. Endpoint gateways enable communications from virtual server instances within your VPC and IBM Cloud service on the private backbone. VPE for VPC gives you the experience of controlling all the private addressing within your cloud.

The VPE supported services page list all of the IBM Cloud services that support VPE and provide links that describe the private hosts to use and any special instructions that might be needed. Not only does this list include the Financial Services Validated IaaS and PaaS services that are available, but they also include a number of platform services, such as Cloud Identity and Access Management (IAM).

Private endpoints should be used whether you are accessing a service by using the CLI, API, or Terraform. For more information, see:

- Securing your connection when using the IBM Cloud CLI
- `visibility` input parameter in Configuring the IBM Cloud Provider plug-in

Creating endpoint gateways

- Create endpoint gateways for the management and workload cluster. For more information, see Creating an endpoint gateway.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-20 Use of External Systems
Security Assessment and Authorization (CA)	CA-3 Information Exchange

System and Communications Protection (SC)	SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
Security Assessment and Authorization (CA)	CA-3 System Interconnections
System and Communications Protection (SC)	SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Connectivity to public internet

Accessing the public internet

The IBM Cloud Framework for Financial Services does not recommend connecting to hosts on the public internet nor accepting connections from the public internet. When it is necessary to do so, you need to use either public gateways or floating IP addresses.

Public gateways and floating IP addresses

A public gateway enables a subnet and all its attached virtual server instances to connect to the internet. Subnets are private by default. After a subnet is attached to the public gateway, all instances in that subnet can connect to the internet. Although each zone has only one public gateway, the public gateway can be attached to multiple subnets.

Floating IP addresses are IP addresses that are provided by the system and are reachable from the public internet.

If you are an application provider from a technology vendor interested in becoming IBM Cloud for Financial Services Validated and want to enable outbound connectivity to the internet, then you will need to make sure you document the data flows and provide evidence that shows how any consumer data that might be leaving the boundary is secured. Similarly, for inbound connectivity, you will need to provide evidence demonstrating how your infrastructure cannot be compromised.

In cases where the IP address of the external service on the internet is not fixed and hence cannot be defined, suggestion is to use a proxy enabled with domain based filtering. Such a proxy needs to reside in an isolated network segment and will be the only component.

See External connectivity for more details on public gateways and floating IP addresses.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-20 Use of External Systems
Security Assessment and Authorization (CA)	CA-3 Information Exchange
System and Communications Protection (SC)	SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
Security Assessment and Authorization (CA)	CA-3 System Interconnections
System and Communications Protection (SC)	SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Working with virtual servers in VPC reference architecture

FQDN based egress proxy in IBM Cloud VPC

Problem statement

FS controls require SG and ACL to not have 0.0.0.0/0 for egress rules. However, some services on internet do not provide specific IP ranges like docker or google, and which are used by ISVs for several purposes from container registry to federation with authentication provider. In high security environments, having an egress for 0.0.0.0/0 is not acceptable. The solution architecture must provide a way to access public services on the internet which do not publish IP Address from isolated secured environments on cloud and still be able to meet FS controls with minimal exceptions.

Assumptions

- At least 1 subnet with bidirectional ACL of 0.0.0.0/0 will be allowed
- At least 1 SG with outgoing rule of 0.0.0.0/0 will be allowed
- Acceptance of increased cost with all internet traffic traversing across Transit GW instead of directly from public GW.
- PGWs can continue to be used in each VPC subnet where the Ips are known and can by specified in ACL/SGs
- Default routing is managed by VPC internally.
- Custom resolver is managed by Private DNS offering

Scope and goals

- Self managed solution
- Sample terraform / automation to set it up.
- HAProxy has been used to implement TCP passthrough proxy
- This guide is only focusing on technical feasibility of an approach. HA considerations must be taken care of on per application basis.

Workload to Internet Connectivity – The BYO Proxy Approach for Services with dynamic IPs

- Deploy a dedicated subnet for hosting the proxy in the Edge VPC. Allow 0.0.0.0/0 on the associated ACL. Only the outgoing SG will have to be configured with 0.0.0.0/0:443 rule.
- The proxy must be configured for whitelisting allowed domains approved by the organization (example *.docker.com)
- Access to services that do not have permanent IP ranges, must use the proxy.
- The proxy can be used as a https proxy.
- This will be marked by SCC and will have to be accepted as an exception.
- For more details refer to public internet access guide

BYO proxy solution overivew

Alternatives to BYO proxy

- TCP Passthrough Load Balancer with domain-based Allow list using SNI headers + Private DNS
  - Does not require any application customization or operating system settings change or networking change in the VPC.
  - This is the recommended approach and the focus of this guide using HAproxy

- Squid (or F5) as a forward HTTP / Socks proxy with domain-based ACL
  - Needs application to support proxy setup or being able to configure the proxy at operating system settings.
- Squid (or F5) transparent proxy as a Default GW with domain-based ACL
  - Needs all the internal VPCs to still have subnets with bidirectional ACLs to allow 0.0.0.0/0(even though no route to internet exists) which will get flagged by SCC. While secure, it is not FS compliant and will fail SCC checks.

High level architecture

High level architecture

All requests to public service endpoints without a DNS alias fail, since there is no public gateway in Management or Workload VPC and hence no route to Internet. Only way to reach the internet is through Edge VPC's TCP Passthrough Proxy, which implements an FQDN allow-list.

Note: Client-to-site VPN is included in the diagram for completeness but is not required for the proxy to function as a firewall. Its primary purpose is to enable access to the isolated environment. Additionally, in accordance with FS controls, the VPN must operate as a full-tunnel VPN. As a result, the proxy will also facilitate controlled internet access for operators connected via VPN, when necessary.

HAproxy configuration example

```
$ # Capture SNI for logging and routing
tcp-request inspect-delay 5s
tcp-request content accept if {{ req_ssl_hello_type 1 }}

# ACL to extract SNI hostname
acl valid_sni req.ssl_sni -i -f /etc/haproxy/allowed_hosts.txt
acl has_sni req.ssl_sni -m found

# Capture SNI for logging
tcp-request content capture req.ssl_sni len 64

# Block connections without SNI or not in ACL
tcp-request content reject unless has_sni
tcp-request content reject unless valid_sni

# Route based on SNI to dynamic backend
use_backend %[req.ssl_sni,lower,map(/etc/haproxy/sni_backend.map,txt,default_backend)]
```

Private DNS configuration

- DNS Services instance must be configured in at least one VPC within the cloud account to provide DNS resolution overrides for the target domains.
- Create a DNS zone in the DNS Services instance for each top level domain that needs to be accessed via the egress proxy. Add DNS ”A” records for each FQDN pointing to the proxy instance.
- For each DNS zone configure permitted subnets that will access the external resources via egress proxy. Do not add the subnet(s) where the proxy VSI resides.
- ICR access: if a VPE for ICR is provisioned within a VPC, the default VPC DNS resolver already points public ICR FQDN to the IBM Cloud private IP addresses, so there is not need to override or proxy connections to the ICR endpoints.
- Restricted DNS zones: certain domain zones are not allowed in DNS Services for security and/or platform stability reasons. In most cases egress access to resources with the FQDNs in these zones can be arranged in a different way - e.g. container image mirroring for image registries such as quay.io or registry.redhat.io, or using forwarding proxy mode for applications that support proxy configuration. In case where none of the alternative access options are feasible, a custom resolver has to be configured in the DNS Services for each subnet that needs access to the resource with a restricted FQDN. The custom resolver should have a matching rule for the restricted FQDN and forward the query to a custom DNS instance (e.g. a BIND name service running on a VSI) that will resolve the FQDN to the egress proxy VSI IP address.

HA considerations

- Instance redundancy
  - Deploy additional instances in other zone to eliminate single points of failure and ensure continued service availability in case of a node or host failure.
- Load balancing
  - While DNS round robin can load balance in case of multi-instance deployment but will not provide health check capabilities. Configure a Load Balancer as a Service (LBaaS) in front of the instances to distribute traffic evenly and provide automatic failover capabilities.
- Scalability and performance
  - The load balancer configuration should support horizontal scaling as needed to handle varying workloads or peak usage scenarios.

- Alignment with Availability Objectives
  - The number of instances and load balancing setup should be determined based on defined RTO (Recovery Time Objective), RPO (Recovery Point Objective), and uptime requirements.

HA Deployment with VPC Application Load Balancers

High availability configuration for BYO proxy deployment

- Follow typical deployment with VPC Application Load Balancers while configuring the VSIs running the proxy component as ALB backend pools.
- Use TCP front-end listener for transparent proxy mode scenario. The private DNS should have CNAME entries for the target external domain pointing to the ALB hostname.
- For transparent proxy listener, configure policies matching the allowed SNI / FQDN - that will give an additional level of control and auditability. Requests to the egress proxy that do not match the target FQDN will result in an SSL error if the default pool is removed. This may cause an SCC ALB rule check failure, so either a dummy backend must be created, or an SCC exception should be added for the ALB handling the proxy traffic.
- For a normal forward proxy mode, it is recommended to set up HTTPS front-end listener with a valid certificate so that the connection to the proxy is properly encrypted. An HTTPS listener can also have policies defined based on HTTP headers, so a destination for a proxy request can be checked as well.
- Using front-end listener policies is not required for make ALB/HA work but it provides an additional layer of security with more visibility at the platform level - as opposed to all configuration encapsulated at the proxy VSI config in its filesystem.
- Define security groups that limit access to the ALB from designated subnets or resources within the VPC. A different security group should restrict traffic to proxy VSI to block direct access to it besides ALB.

Networking configuration examples

Access control lists

Edge internet proxy subnet ACL (edge-internet-proxy-acl)

Inbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
inbound-tcp	1	Allow	TCP	Any IP, any port	10.0.0.0/8, ports 1024-65535	—
inbound-dns	2	Allow	UDP	161.26.0.0/16, port 53	10.0.0.0/8, any port	—
inbound-vpc-to-vpc	3	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Edge internet proxy subnet ACL - inbound rules

Outbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
outbound-tcp	1	Allow	TCP	10.0.0.0/8, ports 1024-65535	Any IP, port 443	—
outbound-dns	2	Allow	UDP	10.0.0.0/8, any port	161.26.0.0/16, port 53	—
outbound-vpc-to-vpc	3	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Edge internet proxy subnet ACL - outbound rules

Management subnet ACL (mgmt-acl)

Inbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
inbound-iaas	1	Allow	ICMP-TCP-UDP	161.26.0.0/16	10.0.0.0/8	—

inbound-cse	2	Allow	ICMP-TCP-UDP	166.8.0.0/14	10.0.0.0/8	—
inbound-vpc	3	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Management subnet ACL - inbound rules

Outbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
outbound-iaas	1	Allow	ICMP-TCP-UDP	10.0.0.0/8	161.26.0.0/16	—
outbound-cse	2	Allow	ICMP-TCP-UDP	10.0.0.0/8	166.8.0.0/14	—
outbound-vpc	3	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Management subnet ACL - outbound rules

Client-to-site VPN subnet ACL (edge-vpn-acl)

Inbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
inbound-internet	1	Allow	UDP	Any IP, any port	10.0.0.0/8, port 443	—
inbound-iaas	2	Allow	ICMP-TCP-UDP	161.26.0.0/16	10.0.0.0/8	—
inbound-cse	3	Allow	ICMP-TCP-UDP	166.8.0.0/14	10.0.0.0/8	—
inbound-vpc	4	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Client-to-site VPN subnet ACL - inbound rules

Outbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
outbound-internet	1	Allow	UDP	10.0.0.0/8, port 443	Any IP, any port	—
outbound-iaas	2	Allow	ICMP-TCP-UDP	10.0.0.0/8	161.26.0.0/16	—
outbound-cse	3	Allow	ICMP-TCP-UDP	10.0.0.0/8	166.8.0.0/14	—
outbound-vpc	4	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Client-to-site VPN subnet ACL - outbound rules

Workload subnet ACL (wrkld-acl)

Inbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
inbound-iaas	1	Allow	ICMP-TCP-UDP	161.26.0.0/16	10.0.0.0/8	—
inbound-cse	2	Allow	ICMP-TCP-UDP	166.8.0.0/14	10.0.0.0/8	—
inbound-vpc	3	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Workload subnet ACL - inbound rules

Outbound rules

Name	Rule Priority	Allow or Deny	Protocol	Source	Destination	Value
outbound-iaas	1	Allow	ICMP-TCP-UDP	10.0.0.0/8	161.26.0.0/16	—

outbound-cse	2	Allow	ICMP-TCP-UDP	10.0.0.0/8	166.8.0.0/14	—
outbound-vpc	3	Allow	ICMP-TCP-UDP	10.0.0.0/8	10.0.0.0/8	—

Workload subnet ACL - outbound rules

Security groups

Egress Proxy VSI Security Group (egress-proxy-sg)

Inbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
TCP	CIDR block	10.0.0.0/8	CIDR block	10.0.0.0/8	Port 443
TCP	CIDR block	10.0.0.0/8	CIDR block	10.0.0.0/8	Port 22

Egress Proxy VSI Security Group - inbound rules

Outbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	166.8.0.0/14	—
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	161.26.0.0/16	—
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	161.26.0.0/16	—
TCP	CIDR block	10.0.0.0/8	Any	0.0.0.0/0	Port 443
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	166.8.0.0/14	—

Egress Proxy VSI Security Group - outbound rules

Management resources Security Group (mgmt-sg)

Inbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	10.0.0.0/8	—

Management resources Security Group - inbound rules

Outbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	161.26.0.0/16	—
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	10.0.0.0/8	—
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	166.8.0.0/14	—

Management resources Security Group - outbound rules

Client-to-site VPN Security Group (c2svpn-sg)

Inbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
UDP	Any	0.0.0.0/0	CIDR block	10.0.0.0/8	Port 443

Client-to-site VPN Security Group - inbound rules

Outbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	10.0.0.0/8	—

Client-to-site VPN Security Group - outbound rules

Workload resources security group (wrkld-sg)

Inbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	10.0.0.0/8	—

Workload resources security group - inbound rules

Outbound rules

Protocol	Source Type	Source	Destination Type	Destination	Value
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	10.0.0.0/8	—
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	166.8.0.0/14	—
ICMP-TCP-UDP	CIDR block	10.0.0.0/8	CIDR block	161.26.0.0/16	—

Workload resources security group - outbound rules

Next steps

- Working with virtual servers in VPC reference architecture

Connecting on-prem environment with service providers on IBM Cloud with Direct Link

Direct Link can provide a secure and dedicated netowork connection between on-prem environment ("Bank") and ISV Cloud account for the ISV to access on-prem services (Bank information systems).

Assumptions

- ISV uses 1 account for multiple workload environments
- VPCs are used for isolation between workload environments
- Internet connectivity to and from the ISV VPCs still follows the regular pattern using Public GW and CIS
- Bank provides the right access to onprem services.
- Bank willing to provides the IP address for the VPC connected to Direct Link
- ISV and the Bank are two distinct organizations, with isolated networking boundaries.

Alternatives

- One Direct Link per VPC (1 per VPC - Management, Workload)
- One Direct Link per provider connected through Transit Gateway to multiple VPCs(Edge, Management, Workload)
- **Recommended**: One Direct Link per provider → Transit VPC (with Proxy/Firewall) managed by provider → Transit Gateway → Multiple VPCs(Different environment Access)
- Overall one Direct Link → Transit VPC (with Proxy/Firewall) managed by consumer → Transit Gateway → Multiple VPCs (Several ISV and their Environment Access)
- One Direct Link per provider → Transit VPC (with PPNLB) managed by provider → Transit Gateway → Multiple VPCs (Different environment Access)
- Overall one Direct Link → Transit VPC (with PPNLB) managed by consumer → Transit Gateway → Multiple VPCs(Several ISV and their Environment Access)

One Direct Link per VPC/environment

One Direct Link per each Environment or VPC - for example, Management, Workload, etc.

One Direct Link per each Environment or VPC

- **Advantages**:
  - **Isolation**: Strong security isolation between VPCs.
  - **No interdependence**: Each link is independent; failures don't affect other VPCs.
  - **Lower intra-VPC complexity**: No need for routing across VPCs.
- **Disadvantages**:
  - **Expensive**: Multiple Direct Link connections = higher port, hosting, and operational costs.
  - **Scalability challenge**: Hard to scale with more VPCs.
  - **Management complexity**: More effort required to manage multiple DLs, FW, SGs, ACLs, IP blocks, and routes.
  - **IPAM**: IP blocks need to be provided by the bank if bank wants to publish 0.0.0.0/0 route for Direct Link (in all cases non overlapping IP Blocks will be required to avoid IP conflict for each VPC directly connected to the Direct Link)

One Direct Link per provider connected through Transit Gateway to multiple VPCs

One Direct Link per provider connected through Transit Gateway to multiple VPCs (Edge, Management, Workload)

One Direct Link per provider

- **Advantages**:
  - **Cost-effective**: Only one Direct Link, reducing port & hosting costs.
  - **Scalable**: Easy to add/remove VPCs via Transit Gateway.
  - **Centralized routing**: Simplifies network architecture and operations.
  - **Lower operational overhead**: Managed connectivity using IBM Transit Gateway.
- **Disdvantages**:
  - **Latency variation**: Traffic hops through Transit Gateway may introduce slight latency.
  - **Security considerations**: Shared connectivity may require additional segmentation/policies, specially since this could expose the DL to Edge VPC
  - **IPAM** – IP blocks need to be provided by the bank if bank wants to publish 0.0.0.0/0 route for Direct Link (in all cases non overlapping IP Blocks will be required to avoid IP conflict for each VPC directly connected to the Direct Link)

One Direct Link per provider to Transit VPC managed by provider

One Direct Link per provider → Transit VPC (with Proxy/Firewall) managed by provider → Transit Gateway → Multiple VPCs

One Direct Link per provider to Transit VPC managed by provider

- **Advantages**:
  - **Security layer**: Centralized inspection via proxy/firewall in Transit VPC.
  - **Policy enforcement**: Unified control over outbound/inbound access.
  - **Cost optimization**: One Direct Link reused for all VPCs.
  - **IPAM**: IP Conflict taken care of by NATing / Proxying. Still requires IP block assignment by Bank for Transit VPC if bank wants to publish 0.0.0.0/0 route for Direct Link. In all cases non overlapping IP Blocks will be required to avoid IP conflict for each VPC directly connected to the Direct Link.
- **Disdvantages**:
  - **Performance bottleneck**: Transit VPC may become a chokepoint if not scaled properly.
  - **Single point of failure**: Proxy/FW failure impacts all traffic.
  - **Operational complexity**: Requires maintaining and scaling the services in Transit VPC.
  - **FW Licensing Costs**

Overall one Direct Link to Transit VPC managed by consumer

Overall one Direct Link → Transit VPC (with Proxy/Firewall) managed by consumer → Transit Gateway → Multiple VPCs

Overall one Direct Link to Transit VPC managed by consumer

- **Advantages**:
  - **Security layer**: Centralized inspection via proxy/firewall in Transit VPC.
  - **Policy enforcement**: Unified control over outbound/inbound access.
  - **Cost optimization**: One Direct Link reused for all VPCs.
  - **IPAM**: IP Conflict taken care of by NATing / Proxying. Still requires IP block assignment by Bank for Transit VPC if bank wants to publish 0.0.0.0/0 route for Direct Link. In all cases non overlapping IP Blocks will be required to avoid IP conflict for each VPC directly connected to the Direct Link.
- **Disdvantages**:
  - **Performance bottleneck**: Transit VPC may become a chokepoint if not scaled properly.
  - **Single point of failure**: Proxy/FW failure impacts all traffic.
  - **Increased management burden**: Consumer must manage Transit VPC, Proxy/Firewall, and route propagation. Responsibility of service availability and performance lies with the consumer.
  - **FW Licensing Costs**

One Direct Link per provider to Private Path managed by provider

One Direct Link per provider → Transit VPC (with PPNLB) managed by provider → Transit Gateway → Multiple VPCs

One Direct Link per provider to Private Path managed by provider

- **Advantages**:
  - **Abstracted access to ISV services** : VPCs can use the same "connector" without direct internet exposure.
  - **Security**: ISV services accessed through a controlled environment (Private Path).
  - **Cost-effective**: One Direct Link and centralized access.
  - **IPAM**: IP Conflict taken care of by NATing / Proxying. Still requires IP block assignment by Bank for Transit VPC if bank wants to publish 0.0.0.0/0 route for Direct Link. In all cases non overlapping IP Blocks will be required to avoid IP conflict for each VPC directly connected to the Direct Link.
- **Disdvantages**:
  - **Performance bottleneck**: Transit VPC may become a chokepoint if not scaled properly.

Overall one Direct Link to Private Path managed by consumer

Overall one Direct Link → Transit VPC (with PPNLB) managed by consumer → Transit Gateway → Multiple VPCs(Several ISV and their environment access)

Overall One Direct Link to Private Path managed by consumer

- **Advantages**:
  - **Consumer control**: Full control over Private Path connectivity, routing, and security policies.
  - **Security segmentation**: ISV services accessed via a consumer-trusted middle layer.
  - **Auditing & compliance**: Easier to monitor and enforce consumer-side policies.
  - **IPAM**: IP Conflict taken care of by NATing / Proxying. Still requires IP block assignment by Bank for Transit VPC if bank wants to publish 0.0.0.0/0 route for Direct Link. In all cases non overlapping IP Blocks will be required to avoid IP conflict for each VPC directly connected to the Direct Link.
- **Disdvantages**:
  - **Performance bottleneck**: Transit VPC may become a chokepoint if not scaled properly.
  - **Increased management burden**: Consumer must manage Transit VPC, Private Path, and route propagation. Responsibility of service availability and performance lies with the consumer.
  - **Single point of failure**: Proxy/FW failure impacts all traffic.

Comparison Summary

Architecture	Cost Efficiency	Scalability	Security	Complexity	Latency	Reliability

One Direct Link per VPC (1 per VPC - Mgmt, Wrkld..)	✕ Low	✕ Low	✓ High	✕ High	✓ Low	✓ High
One Direct Link per provider connected through Transit Gateway to multiple VPCs(Edge, Mgmt, Wrkld..)	⚠ Medium	⚠ Medium	⚠ Medium	✓ Low	⚠ Medium	⚠ Medium
One Direct Link per provider → Transit VPC (with Proxy/Firewall) managed by provider → Transit Gateway → Multiple VPCs(Different environment Access)	⚠ Medium	⚠ Medium	✓ High	⚠ Medium	⚠ Medium	⚠ Medium
Overall one Direct Link → Transit VPC (with Proxy/Firewall) managed by consumer → Transit Gateway → Multiple VPCs (Several ISV and their Environment Access) (**Recommended**)	✓ High	✓ High	✓ High	⚠ Medium	⚠ Medium	⚠ Medium
One Direct Link per provider → Transit VPC (with PPNLB) managed by provider → Transit Gateway → Multiple VPCs(Different environment Access)	✓ High	✓ High	✓ High	⚠ Medium	⚠ Medium	⚠ Medium
Overall one Direct Link → Transit VPC (with PPNLB) managed by consumer → Transit Gateway → Multiple VPCs(Several ISV and their Environment Access)	✓ High	✓ High	✓ Very High	✕ High	⚠ Medium	⚠ Medium

Comparison of Direct Link deployment alternatives

FS guidance for accessing workloads from private intranet

- Workloads may need access from private intranets using VPN or Direct Link.
- **VPN Options**: Client-to-Site and Site-to-Site VPNs support NATing.
- **Direct Link Considerations**:
  - Requires careful IP space coordination to avoid conflicts.
  - Both consumer and provider must align on VPC IP ranges.
- **Recommended Setup**: (One Direct Link per provider → Transit VPC (with Proxy/Firewall) managed by provider → Transit Gateway → Multiple VPCs):
  - Connect Direct Link to a Transit VPC.
  - Use a firewall device in the Transit VPC for NATing.
  - Enables secure and conflict-free workload access.
  - For DNS resolution,
    - If the on-prem environment allowes, a DNS forwarding can be setup using IBM Cloud Private DNS Resolver deployed by ISV in the VPC to forward queries about specific FQDNS/domains to the bank internal DNS.
      - Full Direct Access (with security setup at the internal DNS)
      - 2 Stage – DMZ DNS forwarding to on-prem Internal DNS
    - Alternatively, if allowed by on-prem environment, a non-authoritative zone(alias) can be setup using IBM Cloud Private DNS Resolver which resolves to the appropriate bank internal FQDNS (potential issue : Out of sync DNS records)

See also Consumer connectivity to workload VPC

Next steps

- Data loss prevention and perimeter access patterns

Data loss prevention and perimeter access patterns

Learn about the data loss prevention strategies with a focus on perimeter access patterns for Independent Software Vendor (ISV) deployments in IBM Cloud for Financial Services, including security enforcements, architectural decisions, and threat mitigations.

Overview

This guidance defines the use cases and deployment patterns for ISVs to follow when implementing perimeter access controls in IBM Cloud for Financial Services environments. The patterns address various access scenarios including application users, cloud administrators, and cross-account connectivity.

ISV actors and access patterns

The following table describes the different types of actors in an ISV deployment and their access patterns:

Actor	Privilege level	Access path	Components accessed
Application User	Non-privileged user	Internet access to application	User-facing application components only
Application Admin	Privileged user	Internet access to admin console	Administrative application components
Cloud Admin	Privileged user	VPN and bastion host for cloud resources; internet for Cloud console	Cloud resources within assigned permissions
Service ID (Machine-to-machine)	Privileged user	VPN, trusted profiles	Management VPC resources

ISV actor types and access patterns

Use cases

This guidance covers the following perimeter access use cases:

- Application users (non-privileged) accessing applications in the workload VPC
- Cloud administrators (privileged) accessing cloud resources in the ISV cloud account
- Application administrators (privileged) accessing application admin resources for applications running in workload VPC
- Cloud administrators (privileged) accessing the IBM Cloud console
- Connectivity
  - VPC connectivity using Transit Gateway within the same account
  - VPC connectivity using Transit Gateway across different accounts
  - Internet access using Public Gateway
  - Client-to-site VPN connections from ISV to cloud account
  - Site-to-site VPN connections for privileged access
  - Direct Link connectivity

Reference deployment architecture

Reference deployment architecture

The ISV reference deployment architecture consists of three primary VPCs:

Edge VPC

: Hosts application load balancers, bastion hosts, and VPN gateways. Deployed across 2 availability zones in active-active mode.

Management VPC

: Contains management and operational tools. Deployed across 2 availability zones.

Workload VPC

: Hosts the application workloads. Deployed across 2 or 3 availability zones depending on required SLA.

All cloud services are accessed from VPCs using Virtual Private Endpoints (VPEs), which are deployed for each VPC.

Application user access (non-privileged)

Application users access the application through the internet with multiple layers of security enforcement.

Application user access (non-privileged)

Security enforcements

The following security enforcements are in place for application user access:

1. IBM Cloud Internet Services: Provides Layer 3 DDoS protection, DNSSec, IP firewall (ASN/Region/CIDR filtering), and load balancing.
2. VPC network ACLs: Allow IP addresses only from Cloud Internet Services (CIS).
3. Security groups: Allow connections only from Cloud Internet Services (CIS) to F5.
4. F5 BIG-IP: Provides Layer 7 WAF, DDoS protection, TLS termination, and access policies. Has a public floating IP associated with it (behind ACL and security group).
5. Geographic restrictions: IBM Cloud blocks traffic from embargo countries as documented in IBM Cloud Notices.

Threat mitigation

Compromise of F5 to initiate exfiltration

The only public port open on the external interface is 443 (enforced by security groups and ACLs). When security groups or ACLs are changed, events are logged to IBM Cloud Activity Tracker. For more information, see Activity Tracker events for VPC.

Exfiltration from F5

Public IP cannot be used to log in to F5. Operators must use VPN to access the F5 console. VPN connections are logged by Activity Tracker. If the operator has access only to F5, they cannot change ACLs or security groups. F5 audit logs can be sent to SIEM for monitoring.

Architectural decisions

- Use Cloud Internet Services (CIS) and F5 with appropriate ACLs and security groups
- Send audit logs from F5 to SIEM
- Send Activity Tracker events to SIEM

Cloud administrator access (privileged)

Cloud administrators access cloud resources through VPN and bastion hosts with strict security controls.

Cloud administrator access (privileged)

Security enforcements

The following security enforcements are in place for cloud administrator access:

1. VPC network ACLs: Allow IP addresses only from the ISV network.
2. Security groups: Allow connections only from the ISV network.
3. Client-to-site VPN: Has a public IP but is protected by ACLs and security groups. Requires IAM user credentials and/or certificate-based authentication.
4. IAM access control: Users need IAM access for VPN and bastion host (SSH) access.
5. Bastion host: Records all activities performed by privileged users.
6. Geographic restrictions: IBM Cloud blocks traffic from embargo countries.

Threat mitigation

Unauthorized access

Only IAM-authorized users can connect to the VPN. ACLs and security groups allow access only from the ISV's CIDR ranges.

Exfiltration attempts

VPN connections are logged by Activity Tracker. Bastion hosts log all activity and create recordings.

Architectural decisions

- Use appropriate ACLs and security groups to restrict access

- Send audit logs from bastion hosts to SIEM
- Send Activity Tracker events to SIEM

Application administrator access (privileged)

Application administrators access the application's administrative interface through the internet with enhanced security controls.

Application administrator access (privileged)

Security enforcements

The following security enforcements are in place for application administrator access:

1. Cloud Internet Services (CIS): Provides Layer 3 DDoS protection, DNSSec, IP firewall (ASN/Region/CIDR filtering), and load balancing.
2. VPC network ACLs: Allow IP addresses only from Cloud Internet Services (CIS).
3. Security groups: Allow connections only from Cloud Internet Services (CIS) to F5.
4. F5 public IP: network connections restricted by ACL and security groups
5. F5 capabilities: L7 WAF, DDoS, TLS termination, Access Policy
6. F5 access policies: Can determine if the accessed page is an admin page and apply appropriate policies.
7. Geographic restrictions: IBM Cloud blocks traffic from embargo countries.

Threat mitigation

Unauthorized access

: Allow privileged users to access admin functions only through a bastion host when using VPN access.

Architectural decisions

- Forward application audit logs to SIEM
- If the application has a dedicated URL for admin access, consider routing admin traffic through VPN instead of F5
- If the application does not have a dedicated URL for admin access, use role-based access control (RBAC) at the application level

Cloud console access

Cloud administrators access the IBM Cloud console through the internet with multiple security layers.

Cloud console access

Security enforcements

The following security enforcements are in place for console access:

Account settings

: Allow access only from ISV network IP addresses.

IAM access control

: Provides specific access to appropriate access groups (operators). Individual access policies are not used.

Multi-factor authentication

: 2FA is enforced for all users.

Access groups

: All permissions are managed through access groups, not individual policies.

Activity Tracker

: Logs all changes to resources and provides audit logging for HTTPS access and API calls.

Separation of duties

: Platform access and service access can be separated.

Geographic restrictions

: IBM Cloud blocks traffic from embargo countries.

Threat mitigation

Rogue user from ISV

: Must originate from ISV IP range, have IAM credentials, pass 2FA, have IAM platform permissions, and have IAM service permissions. All actions generate Activity Tracker events for SIEM alerting.

Remote or distributed employees

: Must use full tunnel VPN into ISV enterprise network to access the cloud console. Non-ISV IP range access is denied. Administrators can also access through client-to-site VPN to bastion host, then to cloud console.

Architectural decisions

- Maintain an allow list of IP addresses from which the console can be accessed
- Configure services with the same IP restrictions using Context-based restrictions (CBR)
- Administrators can access from:
  - ISV internal network
  - Client-to-site VPN connection, then to IBM Cloud console
- Add customer IP addresses to allow list for customer users to view tickets on the console

Transit Gateway connectivity (same account)

Transit Gateway enables connectivity between VPCs within the same account with security controls.

Transit Gateway connectivity (same account)

Security enforcements

The following security enforcements are in place for Transit Gateway connectivity within the same account:

1. Prefix filtering: Use Transit Gateway prefix filtering to limit which address prefixes are shared between VPCs.
2. Activity Tracker: Logs all changes to resources and provides audit logging.
3. Network segmentation: Isolate sensitive data and limit the impact of potential breaches.
4. Access controls: Use security groups, network ACLs, and IAM policies to restrict access to sensitive data.
5. Encryption: Ensure data in transit is encrypted.
6. Local Transit Gateway: Use local Transit Gateway for regional deployments.
7. Flow logs analysis: Analyze flow logs using SIEM to ensure source and destination IP addresses are from known CIDR ranges.

Threat mitigation

Inadvertent data sharing

: Sensitive data might be inadvertently shared across VPCs due to misconfigured Transit Gateway connections. Use prefix filtering and network segmentation to prevent this.

Unauthorized data access

: Unauthorized applications or services might leverage Transit Gateway to access and exfiltrate data. Use security groups, ACLs, and flow log analysis to detect and prevent unauthorized access.

Transit Gateway connectivity (cross-account)

Transit Gateway enables connectivity between VPCs across different accounts with additional security controls.

Transit Gateway connectivity (same account)

Security enforcements

The following security enforcements are in place for cross-account Transit Gateway connectivity:

1. Connection approval: Cross-account Transit Gateway connections must be requested and accepted by both accounts.
2. Prefix filtering: Use Transit Gateway prefix filtering to limit which address prefixes are shared between accounts.
3. Activity Tracker: Logs all changes to resources and provides audit logging. For more information, see Activity Tracker events for Transit Gateway.
4. Network segmentation: Isolate sensitive data and limit the impact of potential breaches.
5. Access controls: Use security groups, network ACLs, and IAM policies to restrict access to sensitive data.
6. Encryption: Ensure data in transit is encrypted.
7. Local Transit Gateway: Use local Transit Gateway for tregional deployments.
8. Flow logs analysis: Analyze flow logs using SIEM to ensure source and destination IP addresses are from known CIDR ranges.

Architectural decisions

- Use of cross-account connectivity using Transit Gateway is discouraged and requires addictional justifications
- Enabling cross-account Transit Gateway connectivity can be tracked by Activity Tracker

Threat mitigation

Inadvertent data sharing

Sensitive data might be inadvertently shared across VPCs in different accounts due to misconfigured Transit Gateway connections. Use prefix filtering, connection approval process, and network segmentation to prevent this.

Unauthorized data access

Unauthorized applications or services might leverage Transit Gateway to access and exfiltrate data. Use security groups, ACLs, and flow log analysis to detect and prevent unauthorized access.

Public Gateway access

Public Gateway enables outbound internet connectivity from VPC resources with security controls.

Public Gateway access

Security enforcements

The following security enforcements are in place for Public Gateway access:

1. Access controls: Use security groups, network ACLs, and IAM policies to restrict access.
2. Activity Tracker: Logs all changes to resources and provides audit logging.
3. Encryption in transit: Only allow outbound connections with encryption in transit when outbound connections are needed.
4. Alternative approaches: For sites that do not publish CIDR ranges, assets should be made available using an alternative approach.
5. Security and Compliance Center: Checks for rules that do not allow access to 0.0.0.0/0 with any port, ensuring specific IP addresses or CIDR ranges are used.
6. Geographic restrictions: IBM Cloud blocks traffic from embargo countries.

Threat mitigation

Insider data exfiltration

Insiders with access to instances connected to the internet can intentionally or unintentionally transfer sensitive data outside the organization. Use security groups, ACLs, and monitoring to detect and prevent unauthorized data transfers.

Unauthorized data transfer

Unauthorized applications or services might leverage the Public Gateway to send data to external destinations without proper oversight. Use flow log analysis and SIEM monitoring to detect unauthorized transfers.

Architectural decisions

- Consider using an egress proxy to filter based on URL and domain
- Minimize the number of subnets with access to Public Gateway to mitigate risk

Client-to-site VPN access

Client-to-site VPN enables ISV administrators to connect to cloud resources from remote locations.

Client-to-site VPN access

Security enforcements

The following security enforcements are in place for client-to-site VPN access:

1. Network ACLs: Limit client connections to ISV on-premises networks and Cloud Internet Services (CIS).
2. Cloud Internet Services (CIS) routing: For remote locations, route the client-to-site VPN through Cloud Internet Services (CIS) and use IP firewall to filter by IP, ASN, or region.
3. Access controls: Ensure correct ACLs and security groups are applied to allow access to resources for VPN users.
4. Geographic restrictions: IBM Cloud blocks traffic from embargo countries.

Threat mitigation

Insider data misuse

Employees with legitimate VPN access might misuse their access to transfer sensitive data outside the organization. Use monitoring, logging, and access controls to detect and prevent misuse.

Remote work risks

Remote work environments might lack adequate supervision, increasing the risk of data misuse. Implement endpoint security controls and monitoring.

Architectural decisions

- All internet-exposed interfaces must be restricted to Cloud Internet Services (CIS), with exceptions for site-to-site VPN
- Consider an assessment check for laptop encryption to secure critical data like VPN certificates

Site-to-site VPN access (ISV to cloud)

Site-to-site VPN enables secure connectivity between ISV on-premises networks and cloud resources.

Site-to-site VPN access (ISV to cloud)

Security enforcements

The following security enforcements are in place for site-to-site VPN access:

1. Peer IP restrictions: Site-to-site connections are limited based on peer IP addresses.

2. Activity Tracker: Changes to the connection resource can be tracked by Activity Tracker.
3. Network routing: Remote users connect to on-premises ISV VPN and then can connect to cloud VPC resources with appropriate routing.

Threat mitigation

Insider data misuse

Employees with legitimate VPN access might misuse their access to transfer sensitive data outside the organization. Use monitoring, logging, and access controls to detect and prevent misuse.

Remote work risks

Remote work environments might lack adequate supervision, increasing the risk of data misuse. Implement endpoint security controls and monitoring.

Architectural decisions

- During Financial Services assessment, ISV must demonstrate how they handle security restrictions to minimize access to IBM Cloud account
- Define controls for using VPN for machine-to-machine flows

Site-to-site VPN access (service consumer to cloud)

Site-to-site VPN enables secure connectivity between service consumer on-premises networks and cloud resources.

Site-to-site VPN access (service consumer to cloud)

Security enforcements

The following security enforcements are in place for customer site-to-site VPN access:

1. Peer IP restrictions: Site-to-site connections are limited based on peer IP addresses.
2. Activity Tracker: Changes to the connection resource can be tracked by Activity Tracker.
3. Network routing: Remote users connect to on-premises customer VPN and then can connect to cloud VPC resources with appropriate routing.

Threat mitigation

Insider data misuse

Employees with legitimate VPN access might misuse their access to transfer sensitive data outside the organization. Use monitoring, logging, and access controls to detect and prevent misuse.

Remote work risks

Remote work environments might lack adequate supervision, increasing the risk of data misuse. Implement endpoint security controls and monitoring.

Direct Link connectivity

Direct Link provides dedicated, private connectivity between on-premises networks and IBM Cloud.

Direct Link connectivity

Security enforcements

The following security enforcements are in place for Direct Link connectivity:

1. Peer IP restrictions: Direct Link connections are limited based on peer IP addresses configured at the point of presence (PoP).
2. Activity Tracker: Changes to the connection resource can be tracked by Activity Tracker.
3. Network routing: Remote users connect to on-premises ISV VPN and then can connect to cloud VPC resources with appropriate routing.

Threat mitigation

Insider data misuse

: Employees with legitimate Direct Link access might misuse their access to transfer sensitive data outside the organization. Use monitoring, logging, and access controls to detect and prevent misuse.

Remote work risks

: Remote work environments might lack adequate supervision, increasing the risk of data misuse. Implement endpoint security controls and monitoring.

Architectural decisions

- Similar to site-to-site VPN, define how the access through Direct Link is restricted
- Establish audit procedures to verify access controls

ISV - Application Deployment

ISV - Application Deployment

High availability considerations

The following high availability patterns are recommended for ISV deployments:

Workload VPC

: Deploy across 2 or 3 availability zones depending on required SLA.

Management VPC

: Deploy across 2 availability zones.

Edge VPC

: Deploy across 2 availability zones. Public application load balancers are deployed in active-active mode. Optionally, use separate Edge VPCs for critical and non-critical workloads. The Edge VPC hosts application load balancers, bastion hosts, and VPN gateways.

Cloud Internet Services (CIS)

: Deploy WAF and global load balancer with Cloud Internet Services (CIS).

Cloud service access

All IBM Cloud services are accessed from VPCs using Virtual Private Endpoints (VPEs). VPEs are deployed for each VPC to ensure private connectivity to cloud services.

Next steps

- Working with virtual servers in VPC reference architecture

Networking for Satellite reference architecture

Satellite networking overview

Properly setting up and managing your networking components is critical to building a solution that meets the best practices and requirements of the IBM Cloud Framework for Financial Services. You are at the starting point for learning how to complete the tasks for the Satellite reference architecture.

Next steps

Walk through the topics for how to build out your networking solution:

- Ensuring isolation between Satellite management functions and workload functions
- Accessing external resources from the Satellite location

- Workload operator actions through a bastion host
- Connectivity to IBM Cloud services with Satellite Link
- Consumer connectivity to workload resources

Ensuring isolation between Satellite management functions and workload functions

A key aspect of the IBM Cloud Framework for Financial Services is to separate user workloads from system management functions and isolate security functions from nonsecurity functions. The network infrastructure of the Satellite location can be used to provide physical and logical separation between the Satellite management control plane and your workloads.

Network flow rule design should follow the IBM Cloud Framework for Financial Services's information flow guidelines by using a "deny by default" approach.

Before you begin

1. Complete the work for account setup and management.
2. Complete Satellite location setup.

Identify network areas for control plane hosts and workload hosts

1. Place control plane hosts into a separate network segment. Control plane hosts support various management and security-related components of the Satellite location. To facilitate effective network flow restrictions within the Satellite location, it is recommended to place control plane hosts into a separate network segment (whether physical or virtual) that can enable clear identification of source or destination of the network flows related to control plane functionality. The control plane hosts can be deployed to different physical locations, but the address space they are assigned to should provide an easy way to identify this group of hosts (for example, CIDR blocks).
2. Designate a separate network segment for each group of Satellite hosts assigned to Red Hat OpenShift on IBM Cloud workload clusters. Satellite hosts that are assigned to Red Hat OpenShift on IBM Cloud workload clusters (workload hosts) should use their own network segments that would enable network flow control and monitoring between workload hosts, control plane, and other components outside of the Satellite location. It is recommended to designate a separate network segment (virtual subnet, VLAN, or a similar entity) for each group of Satellite hosts assigned to the same workload cluster.
3. Ensure network flows between network segments that are allocated for control plane hosts and workload hosts are allowed without restrictions. Therefore, in general, the hosts within the control plane or same workload cluster need to have unrestricted network connections to each other, the network flow control is applied to traffic that crosses the boundaries of a particular group of hosts that are assigned to the control plane or same workload cluster. The host network requirements provide detailed specifications for connectivity requirements between Satellite location hosts.
4. Make sure that the selected network infrastructure can provide performance compatible with the host latency requirements.

Design an addressing plan for control and workload segments

1. Define network address range (subnet or CIDR block) for control plane hosts.
2. Define network address range (subnet or CIDR block) for the workload hosts of each Red Hat OpenShift on IBM Cloud cluster in the Satellite location.

Configure network flow rules to restrict network traffic for control plane

The following rules for control plane hosts must be implemented within the networking infrastructure (virtual or physical) provided by the operator of the Satellite location.

1. Allow inbound flows for control plane hosts (API endpoints, TCP 30000 - 32767) from management plane within your environment. For example, to support bastion or CI/CD tools.
2. Allow bidirectional flows between control plane hosts and workload hosts of the same Satellite location.
3. Allow outbound flows to workload hosts of the same Satellite location.
4. If Satellite Link location endpoint is configured to enable connection from IBM Cloud to a resource within the Satellite location, allow the outbound flow from control plane hosts to the resource used in the location endpoint for each such endpoint.
5. If Satellite Link cloud endpoints are used by components other than workload hosts in the Satellite location (for example, to access IBM Cloud IAM services from an application deployed outside of the Satellite location), allow inbound flow to the control plane hosts from the resource that requires access the cloud endpoint that uses the destination port listed in the endpoint address attribute.
6. Ensure that control plane hosts are configured to allow outbound access to IBM Cloud public endpoints. For more information, see Accessing external resources from the Satellite location.

Configure network flow rules for workload hosts

The following rules for workload hosts must be implemented within the networking infrastructure (virtual or physical) provided by the operator of the Satellite location.

1. Allow inbound network flows to workload hosts from management resources within your environment, such as a bastion host in your management plane. It is recommended to use networking components capable of controlling Level 7 traffic to limit access to services such as Red Hat OpenShift on IBM Cloud web console based on the FQDN of the request. A combination of network rules and an HTTP proxy can be used as well.
2. Allow bidirectional network flows between workload hosts and control plane hosts of the same Satellite location.
3. Identify and allow network flows to other application resources within the environment that is required by the deployed workloads. For example, databases that are not deployed on the same workload service.
4. Identify and allow network flows to other services deployed in the Satellite location (for example, another Red Hat OpenShift on IBM Cloud cluster) if required by the workload application architecture.
5. Identify and allow the ingress network flows from the edge plane that is used to expose the workload services to application consumers (for example, load balancers).
6. Allow outbound network flows to service endpoints in management plane that is needed for logging and monitoring.
7. Ensure that workload hosts have outbound access to IBM Cloud public endpoints. For more information, see Accessing external resources.

(Optional) Configure virtual network flow rules within Red Hat OpenShift on IBM Cloud

1. Configure virtual network flow rules within Red Hat OpenShift on IBM Cloud. In addition to the network flow controls implemented within the IaaS layer of the Satellite location, you can use virtual networking features of Red Hat OpenShift on IBM Cloud to control network flows within and into your cluster. You can use Kubernetes network policies or Calico network policies to define network flow restrictions for each workload deployed on your Red Hat OpenShift on IBM Cloud cluster.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls below are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-20 Use of External Systems
Security Assessment and Authorization (CA)	CA-3 Information Exchange
System and Communications Protection (SC)	SC-2 Separation of System and User Functionality SC-3 Security Function Isolation SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
Security Assessment and Authorization (CA)	CA-3 System Interconnections
System and Communications Protection (SC)	SC-2 Application Partitioning SC-3 Security Function Isolation SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Accessing external resources from the Satellite location

Accessing external resources from the Satellite location

In general, the IBM Cloud Framework for Financial Services does not recommend connecting to hosts on the public internet nor accepting connections from the public internet. When it is necessary to do so, you need to use proper facilities, such as load balancers, public gateways, or proxy servers, that usually would be deployed outside of the Satellite location network.

However, deployment of Satellite hosts requires connection to certain external resources to enable the attachment and assignment processes for the hosts as well as ongoing operation of Satellite control plane and workload clusters. These connections need to be properly configured, monitored, and audited.

External resources for control plane hosts

Control plane hosts require the following non-HTTP connectivity:

- Red Hat NTP service (unless configured with local or private NTP pool)
- Control plane master connections (TCP ports 30000 - 32767)

The connectivity to the following endpoints must also be allowed, but can be potentially facilitated through an HTTP or HTTPS proxy for flow control and auditing:

- Satellite Link tunnel server endpoint
- IBM Cloud general APIs and Container services
- IAM REST APIs
- LaunchDarkly service
- Object Storage for etcd backup
- Attachment / assignment endpoint
- Satellite Config and Link APIs
- RHEL container registry
- IBM Cloud Container Registry
- IBM Cloud monitoring and log analysis

For more information, see host networking requirements.

External resources for workload hosts

Workload hosts require the following non-HTTP connectivity:

- Red Hat NTP service (unless configured with local or private NTP pool)

The connectivity to the following endpoints must also be allowed, but can be potentially facilitated through an HTTP or HTTPS proxy for flow control and auditing:

- IBM Cloud general APIs and Container services
- IAM REST APIs
- LaunchDarkly service
- Attachment / assignment endpoint
- Satellite Config and Link APIs
- IBM Cloud Container Registry
- RHEL container registry
- IBM Cloud monitoring and log analysis

For more information, see more host networking requirements

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control

Access Control (AC)	AC-20 Use of External Systems
Security Assessment and Authorization (CA)	CA-3 Information Exchange
System and Communications Protection (SC)	SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
Security Assessment and Authorization (CA)	CA-3 System Interconnections
System and Communications Protection (SC)	SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Completing operator actions through a bastion host

Completing workload operator actions through a bastion host

All interactive operator actions on the resources in Satellite location must be run through a bastion host in the management plane. A bastion host is a server that can be accessed through SSH, Windows Remote Desktop Protocol (RDP), `oc` (Red Hat OpenShift CLI), or `kubectl`, but only through a private secure connection. After set up, the bastion host allows a secure connection to Red Hat OpenShift on IBM Cloud clusters, Satellite APIs, or Satellite location hosts. Administrative tasks on the individual services or hosts are completed by using SSH, `oc`, or `kubectl` proxied through the bastion. In addition, session auditing must be enabled to record all privileged user actions.

After connecting to your management plane through your designated connection method specific to your environment, a complete bastion solution should meet the following requirements.

- Full session recording of actions executed via the bastion to perform session audits. This includes SSH, and Kubernetes exec (`kubectl exec -it`) sessions.
    - Linux-based session recordings are captured as a raw dump of stdout and stderr streams, including TAB characters (bash escape sequences).
    - Session recordings should be placed in to a storage service maintained within your environment for archiving. The storage should be configured as “immutable object storage.” Retention policies are should be applied to the storage, so that data is stored in a WORM (Write-Once-Read-Many), non-erasable and non-rewritable manner. The policy is set and enforced for a 12-month (minimum) retention period.
- Authenticating users with MFA using a physical hardware-based security key that generates a six-digit numerical code. A smart card or hardware token designed and operated to FIPS 140-2 level 2 or above or equivalent (e.g., ANSI X9.24, ISO 13491-1:2007) is recommended.
- Locking user accounts after three (3) consecutive failed logon attempts within 15 minutes.
- Locking user accounts for 30 minutes when there have been more than three unsuccessful logon attempts. After the lockout period ends, the user will be able to reset their password. Internal privileged accounts must remain locked until released by an administrator.
- Session timeout after 15 minutes of inactivity.
- Providing a system use notification banner from either the bastion or the target system. The warning banner is displayed before the system grants access to the user and the usage conditions must be approved before proceeding. The banner will provide privacy and security notices consistent with applicable customer policies, regulations, standards, and guidance. The warning banner will state that:
    - Users are accessing a financial services information system.
    - Information system usage may be monitored, recorded, and subject to audit.
    - Unauthorized use of the information system is prohibited and subject to criminal and civil penalties.
    - Use of the information system indicates consent to monitoring and recording.

Access to Red Hat OpenShift on IBM Cloud management APIs

Red Hat OpenShift on IBM Cloud management tasks can be completed by using `oc` CLI tool and `kubectl`. If your bastion solution provides direct support for `kubectl` or `oc` sessions, it is recommended to configure such capabilities for the Red Hat OpenShift on IBM Cloud clusters that are deployed to your Satellite locations. Otherwise, you should consider setting up hosts in your management plane that can provide shell access with `oc` or `kubectl` CLI tools to complete Red Hat OpenShift on IBM Cloud management. Operators should access those hosts through SSH connecting through a bastion.

Access to Red Hat OpenShift on IBM Cloud management UI

Red Hat OpenShift on IBM Cloud cluster that is deployed on a Satellite location provides a management UI (web console). You should consider UI access as an exception, preferring the operator actions should CLI or CI/CD automation for effective access control, reliability, and audit.

If the operators or administrators need to access this UI, the access should be facilitated by the bastion solution as well. In this scenario, a bastion solution would act as an HTTPS reverse proxy that the operator would use to access the Red Hat OpenShift on IBM Cloud management UI.

Access to Satellite hosts through SSH

In general, the Satellite architecture discourages shell to the host OS, including SSH. Host OS updates and other maintenance (for example, backups) should be automated or completed by creating and attaching a new host instance. Login with a `root` account is explicitly blocked after a host is attached to a Satellite location.

If operator access to a Satellite host is required for troubleshooting purposes, SSH connection should be completed through a bastion that uses a nonroot account.

Set up a bastion host

You need to install and manage your own bastion solution within your management plane. A bastion solution can be implemented in a variety of ways. For one example that uses Teleport Enterprise Edition in the VPC reference architecture, see Setting up a bastion host for secure connectivity.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Audit and Accountability (AU)	AU-14 Session Audit
Access Control (AC)	AC-6 (9) Least Privilege | Log Use of Privileged Functions AC-17 Remote Access
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (organizational Users) | Multi-factor Authentication to Privileged Accounts

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Audit and Accountability (AU)	AU-14 Session Audit
Access Control (AC)	AC-6 (9) Least Privilege | Auditing Use of Privileged Functions AC-17 Remote Access
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (Organizational Users) | Network Access to Privileged Accounts

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Using IBM Cloud services with Satellite Link

Connectivity to IBM Cloud services with Satellite Link

When using IBM Cloud services from workload components in your Satellite location, you should use Satellite Link functionality to facilitate secure access to the IBM Cloud services. Satellite Link endpoints provide a connection over a secure tunnel directly to private service endpoints.

Inside a Satellite location, this private access can be accomplished by using a Satellite Link endpoint to map a TCP port on the control plane hosts to the IBM Cloud service private endpoint. The Link endpoint is a virtualized function that scales horizontally, is redundant and highly available, and spans all

zones of your Satellite location.

> ⚠ **Important:** Private service endpoints for IBM Cloud services should be used when you define cloud endpoints for your location.

Creating Satellite Link endpoints

- Create Satellite Link endpoints for the IBM Cloud services that you need to consume from your Satellite location. For more information, see Creating and managing link endpoints.

> ⚠ **Important:** Cloud endpoints for your location use a TCP port on the control plane hosts. If the endpoint is supposed to be used by a component outside of your Red Hat OpenShift on IBM Cloud cluster deployed on your Satellite location, you need to allow network connectivity from the origin of that component to the port on the control plane.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-20 Use of External Systems
Security Assessment and Authorization (CA)	CA-3 Information Exchange
System and Communications Protection (SC)	SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-20 Use of External Information Systems
Security Assessment and Authorization (CA)	CA-3 System Interconnections
System and Communications Protection (SC)	SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Consumer connectivity to workload resources

Consumer connectivity to workload resources

Consumers of the workloads and services that are deployed to your Satellite location might need access to one or more of the following resources:

- Applications and microservices that are deployed on your Red Hat OpenShift on IBM Cloud cluster
- Red Hat OpenShift on IBM Cloud Link cloud endpoints providing access to IBM Cloud services.

Configure your network infrastructure so that access to workloads and services that are not related to security or management functions is provided through an edge plane that runs outside of your Satellite location or some other facility that can control network flows to your Satellite hosts. Make sure that services and workloads related to security and management activities can be accessed only from the management plane that runs outside of your Satellite location.

Edge plane with web application firewall

The edge plane should be used to enhance boundary protection for the workloads and services that are deployed in the Satellite location and the Satellite control plane. The workload consumer connectivity to the edge plane is your responsibility and depends on the requirements for the workload consumers and the existing connectivity options for your network infrastructure.

> ⚠ **Important:** We recommend that you use a Web Application Firewall (WAF) in your edge plane for any service or workload that uses the HTTP protocol (including REST APIs). A WAF filters and monitors consumer traffic and can prevent attacks that exploit the vulnerabilities of web applications and services.

Exposing your Red Hat OpenShift on IBM Cloud workloads

Red Hat OpenShift on IBM Cloud provides several options for enabling external access to applications deployed on a cluster:

- Ingress controller
- Load balancer service
- External IP
- NodePort

For more information, see Exposing apps in Satellite clusters, About Ingress, and Ingress traffic configuration.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
System and Communications Protection (SC)	SC-5 Denial-of-service Protection SC-7 Boundary Protection SC-7 (4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny by Default - Allow by Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
System and Communications Protection (SC)	SC-5 Denial of Service Protection SC-7 Boundary Protection SC-7(4) Boundary Protection | External Telecommunications Services SC-7 (5) Boundary Protection | Deny By Default - Allow By Exception SC-7 (10) Boundary Protection | Prevent Exfiltration SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-11 Trusted Path

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Working with Red Hat OpenShift on IBM Cloud

Compute and containers

Working with virtual servers in VPC reference architecture

Virtual Servers for VPC is an infrastructure-as-a-service (IaaS) offering that gives you access to all of the benefits of VPC, including network isolation, security, and flexibility. You can quickly provision instances with high network performance. When you provision an instance, you select a profile that matches the amount of memory and compute power that you need for the application that you plan to run on the instance. Instances are available on the x86 architecture.

You can optionally use dedicated hosts for VPC. You can create a dedicated host to carve out a single-tenant compute node, free from users outside of your organization. Within that dedicated space, you can create virtual server instances according to your needs. Additionally, you can create dedicated host groups that contain dedicated hosts for a specific purpose. Because a dedicated host is a single-tenant space, only users within your account that have the required permissions can create instances on the host.

Tip: Dedicated hosts are highly recommended when you use virtual servers -- particularly for any parts of your application that process regulated data and keep it in memory.

Deployment

1. (Optional) Create a dedicated host if you want to carve out a single-tenant compute node for your virtual servers. See Creating dedicated hosts and groups for more details.

   Note: Dedicated hosts are highly recommended, particularly for any parts of your application that process regulated data and keep it in memory.

2. Create one or more virtual servers.
   - If using dedicated hosts, see the following for more details:
     - Creating instances on dedicated hosts
   - If not using dedicated hosts, see the following for more details:
     - Planning for instances

Next steps

- If you plan to use Red Hat OpenShift on IBM Cloud, see Working with Red Hat OpenShift on IBM Cloud.
- Otherwise, see Audit logging of IBM Cloud events.

Working with Red Hat OpenShift on IBM Cloud

If you want to use containers in either either the VPC or Satellite reference architectures, you should use Red Hat OpenShift on IBM Cloud. Red Hat OpenShift on IBM Cloud is a managed offering to create your own Red Hat OpenShift on IBM Cloud cluster of compute hosts to deploy and manage containerized apps on IBM Cloud. Red Hat OpenShift on IBM Cloud provides intelligent scheduling, self-healing, horizontal scaling, service discovery and load balancing, automated rollouts and rollbacks, and secret and configuration management for your apps. Combined with an intuitive user experience, built-in security and isolation, and advanced tools to secure, manage, and monitor your cluster workloads, you can rapidly deliver highly available and secure containerized apps in the public cloud.

Deploying Red Hat OpenShift on IBM Cloud

1. Install the CLI plugins for Red Hat OpenShift on IBM Cloud. For more information, see Installing the Red Hat OpenShift on IBM Cloud CLI.
2. Setup the API for Red Hat OpenShift on IBM Cloud. For more information, see Setting up the API.
3. Create your Red Hat OpenShift on IBM Cloud cluster. For more information, see Creating a Red Hat OpenShift on IBM Cloud cluster in your VPC.
4. Install Red Hat OpenShift on IBM Cloud Service Mesh which is based on the open source Istio project.

   Two of the most important reasons for using Red Hat OpenShift on IBM Cloud Service Mesh is to enable you to:

   - Encrypt network traffic between microservices running in your cluster. See Data encryption in transit and enable mTLS between containers for more details.
   - Implement gateways to specify which traffic you want to enter or leave the mesh (and deny all traffic by default). You can use an egress gateway to control/allowlist all necessary endpoints and domains that your application needs to connect to. For examples, see Expose the app

with the Istio Ingress Gateway and Route and Perform traffic management

5. Set up Container Registry and Vulnerability Advisor. For more information, see Container Registry and Vulnerability Advisor.
6. Develop and deploy applications to your cluster. See the following for more details:
    - Planning app deployments for more details.
    - Service Mesh on Red Hat OpenShift on IBM Cloud
    - Deploying applications on Red Hat OpenShift on IBM Cloud Service Mesh

Related resources

- Security for Red Hat OpenShift on IBM Cloud
- Scalable web application on Red Hat OpenShift on IBM Cloud
- Deploy microservices with Red Hat OpenShift on IBM Cloud

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (organizational Users) | Multi-factor Authentication to Privileged Accounts
System and Communications Protection (SC)	SC-7 (5) Boundary Protection | Deny by Default / Allow by Exception SC-8 Transmission Confidentiality and Integrity

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Identification and Authentication (IA)	IA-2 (1) Identification and Authentication (Organizational Users) | Network Access to Privileged Accounts
System and Communications Protection (SC)	SC-7 (5) Boundary Protection | Deny by Default / Allow by Exception SC-8 Transmission Confidentiality and Integrity

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

If using the VPC reference architecture, see Storage for VPC reference architecture.

If using the Satellite reference architecture, see Storage for Satellite reference architecture.

Container image management

Image management in Red Hat OpenShift on IBM Cloud clusters

This pattern presents a set of architectural alternatives for controlling access to container images used in workload deployments on Red Hat OpenShift on IBM Cloud clusters. The result is an approach that enables secure workload deployment with required controls over software supply chain elements.

IBM Cloud for Financial Services provides several controls applicable to container image management for workload deployment on IBM Cloud container platform services such as Red Hat OpenShift on IBM Cloud :

- SI-7 - Software, Firmware, and Information Integrity
- CM-6 - Configuration Settings
- SI-4 (4) - Inbound and Outbound Communications Traffic
- SC-7 - Boundary Protection
- SC-7 (5) - Deny by Default — Allow by Exception

A typical workload deployment on an Red Hat OpenShift on IBM Cloud cluster usually requires a number of container images often pulled from public

registries. A secure and compliant solution requires implementing control over both the inventory of the allowed software components and the network access to its distribution point such as a container registry.

Scope

This pattern addresses the following aspects of container image management:

Repository access

- Use only authorized managed repositories (such as IBM Container Registry on IBM Cloud) for container images. The image registry service or instance should comply with the requirements for account and identity management, hosting infrastructure security, observability, and resilience.
- Restrict network access to external public registries while enabling deployment of required images using registry mirroring.
- Prevent workload deployment process from pulling images from an arbitrary external registry or not in the approved list to the container platform directly.

Image acceptance policies

- Only specific images from an approved list should be allowed to be deployed on a container platform.
- All applicable software components must be cryptographically signed by the manufacturer or developer to ensure you can perform integrity checks and verify that the component that is deployed in the system is the same component that the manufacturer or developer evaluated and certified.
- Controlled deployment of Operators from public marketplaces using registry mirroring.

Vulnerability scanning

- Container images stored in a managed registry should be periodically scanned for vulnerabilities.
- The images deployed on a container platform or in the approved list should be periodically scanned for vulnerabilities.

Out of scope

The following topics are not covered in this pattern:

- Container platforms other than Red Hat OpenShift on IBM Cloud
- Secure development process for building custom images and deploying cluster resources (CI/CD)
- A procedure to review, validate, and approve container images required for the workload, including common public images
- A procedure to respond to new vulnerabilities discovered in workload images

Approaches and recommendations

Depending on your requirements for external or third-party container images, you can implement various scenarios for image management.

In any of the scenarios, it is recommended to have the required container images stored in a private namespace in Container Registry for full control of the image inventory and network access to the registry.

This pattern outlines the following approaches:

Explicit references to images in Container Registry private namespaces

If you can configure your deployment manifests to reference container images in your private Container Registry, you can avoid a need to access external or public or third-party registries. This option is most suitable for small custom workloads where you can modify the image references in your deployment manifests to point to a private Container Registry namespace.

Explicit references to images in Container Registry private namespaces

Image mirroring to Container Registry for transparent image reference resolution

The registry mirroring can be set up on cluster nodes to use alternative locations for certain registries and images. Instead of pulling an image from an external registry, a cluster node will use a reference pointing to a private Container Registry namespace. Image mirroring can be used in case of more complex workload deployments where it is not practical to modify the image references in the original deployment manifests, for example, when the workloads are deployed by existing Helm charts.

Image mirroring to Container Registry

Operator index and bundle image mirroring to Container Registry for operator deployments

When deploying software components on an Red Hat OpenShift on IBM Cloud cluster via operators, various images are retrieved from one or more external

image registries. Since you cannot control the image references in the original operator bundle images, image mirroring must be used along with mirroring operator catalog index images and operator bundle images to avoid opening network path to the external registries.

In a secure deployment of Red Hat OpenShift on IBM Cloud cluster, the cluster nodes do not have direct access to the external container registries. When a cluster node tries to pull an image referencing an external registry, the registry mirror configuration can redirect that request to a private Container Registry namespace.

A process to mirror the required images to Container Registry needs to be set up in an environment separate from the workload Red Hat OpenShift on IBM Cloud cluster – for example, a management VPC that has an HTTPS Proxy allowing a connection to the external registry or a compute resource on an enterprise network that can reach both the external registry and the Container Registry private endpoint over a secure connection to IBM Cloud.

Next steps

- Review planning guidance for Red Hat OpenShift on IBM Cloud administrators
- Understand key considerations for evaluating image management options
- Follow the implementation guides to set up image management for your cluster

Guidance for Red Hat OpenShift on IBM Cloud administrators

Review the following guidance to plan your container image management strategy for Red Hat Red Hat OpenShift on IBM Cloud clusters on IBM Cloud VPC.

Outbound traffic protection and external registries

The System and Communication Protection controls in IBM Cloud for Financial Services framework (for example, SC-7 (5)) require that outbound connections from an Red Hat OpenShift on IBM Cloud cluster are explicitly permitted by destination. This creates a problem with accessing external image registries which often use a wide range of IP addresses that can change.

While HTTPS traffic can be routed through a transparent or forwarding proxy to enforce only allow-listed destinations, using this approach with an Red Hat OpenShift on IBM Cloud cluster in VPC still does not allow for fine-grain control over the image sets from external registries that are permitted (only the whole registry can be on the allow-list) and can lead to other issues.

Instead, it is recommended to store the container images in a private namespace in IBM Cloud Container Registry (Container Registry) which can be accessed via a Virtual Private Endpoint (VPE) that provides an even greater control over network flows while also significantly reducing traffic volume to a public network.

Container image integrity

Once all the required container images are stored in a private Container Registry namespace, only the images from the private namespace that you control can be deployed on the Red Hat OpenShift on IBM Cloud cluster, since the network path to external registries is blocked. Container Registry also provides the image scanning functionality, therefore you will have full visibility for the vulnerabilities that are discovered in the set of images being used in your workloads.

The Red Hat OpenShift on IBM Cloud cluster will still have access to baseline images on public namespaces in Container Registry that are required for the default deployment of Red Hat OpenShift on IBM Cloud standard components. You can restrict the images being deployed in your own namespaces by using image acceptance policies.

Operator deployment

If your solution involves deploying operators from public or private catalogs, you will need to identify the required operators and their source registries. Red Hat OpenShift on IBM Cloud provides a command line tool (`oc-mirror`) that can automate mirroring of images used to deploy operators into a private Container Registry namespace.

For each of the operators, you will need to go through the steps of mirroring the operator images into your Container Registry private namespace, including catalog index and operator bundle images, as well as properly configuring Catalog Source resources for the mirrored operator index image. After that, the selected set of operators will become available for installation in OpenShift.

Image mirroring

The process that copies the images from the source registry to your private Container Registry namespace (for example, using OpenShift's `oc-mirror` tool) needs to have network access to both the source external registry and the destination private endpoint in Container Registry.

Depending on the location of the source registry, you may need to use a forwarding proxy that permits the HTTPS connections to the external registry endpoint from a compute environment on IBM Cloud. It is recommended to run the mirroring process outside of the Red Hat OpenShift on IBM Cloud cluster to decouple the network permissions for image mirroring from the outbound traffic restriction for the cluster itself.

The following options describe the mirroring process that uses Container Registry private endpoints to ensure controlled access to Container Registry private namespaces where a public endpoint access is blocked by Content Based Restrictions. Direct access to Container Registry public endpoints from

external environment should be considered a deviation from the Financial Services framework guidelines that requires a separate risk assessment and mitigation.

Option 1: Code Engine job to run the mirroring process

IBM Cloud Code Engine can host a job running `oc-mirror` tools when the job execution is requested.

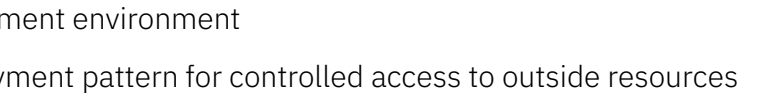

Mirroring to Container Registry with Code Engine application

Advantages:

- Minimal configuration in IBM Cloud
- Secure handling of necessary credentials via Code Engine secrets
- Lower operational and cloud resource cost

Disadvantages:

- Code Engine outbound traffic can be controlled via IP ranges, but access to certain public registries may require all outbound traffic from Code Engine project to be permitted
- An image for `oc-mirror` tool needs to be built

Option 2: Mirroring process on a VSI in Management VPC

The `oc-mirror` CLI tool can be set up on a VSI within a Management VPC. An administrator will log in to the VSI shell and run the command manually when an update to the mirror is required.

Mirroring to Container Registry with a VSI in Management VPC

Advantages:

- Quick to set up within an existing management environment
- Uses common HTTPS egress proxy deployment pattern for controlled access to outside resources

Disadvantages:

- Requires Management VPC and HTTPS proxy configuration
- Managing credentials for Container Registry access may require more complex setup with VSI trusted profiles

Option 3: Mirroring process running within an enterprise environment outside of IBM Cloud

In case the source images are pulled from a private registry available only within an enterprise network, the `oc-mirror` process can be executed on a compute resource (a virtual machine or a containerized environment on-premises) within the private enterprise network.

Pushing the image to the Container Registry is performed via a VPN tunnel or a similar secure connection between the enterprise network and IBM Cloud to use the Container Registry private endpoint. This option can also be used when pulling images from a public external registry if the enterprise network can facilitate a controlled outbound connection to the registry.

Mirroring to Container Registry from an on-prem environment

Advantages:

- Enables mirroring images from registries on private networks
- Minimal configuration in Management VPC

Disadvantages:

- Requires VPN connectivity to Management VPC
- Requires compute facilities within the enterprise network

Image and operator update strategy

Since the images are copied from external registries by `oc-mirror` when the job is executed, the new images or changes in the source image tags (for example, which image is tagged as `latest`) are only appearing in the target Container Registry after the next `oc-mirror` run.

The recommended approach is to explicitly list the tags or versions of operators and images in the `oc-mirror` configuration to avoid unexpected changes in the images and operators being deployed.

Operator updates

When a new version of an operator is made available on the source registry, it should be evaluated in a staging environment for potential adverse upgrade effects and then added to the `oc-mirror` configuration. After the `oc-mirror` run, the new version will become visible in the catalog view on the cluster.

Image updates

For individual images, the specific version tags should be included in `oc-mirror` configuration and used in deployments. Newly released versions should be evaluated and added to the mirroring list.

When referencing dynamic tags such as `latest`, `oc-mirror` will copy the image that currently has the specified tag and will re-assign the tag to the copied image. As a result, the previous image in Container Registry with that tag will lose the tag.

`oc-mirror` does not copy all the tag assignments to the target registry, only assigning the tag explicitly given in configuration. It may be necessary to manually assign additional tags in Container Registry if the deployment process depends on the multiple tags assigned to an image.

Next steps

- Review key considerations for evaluating image management options
- Start with Container Registry configuration

Key considerations for evaluating image management options

Review key considerations for evaluating container image management strategies on IBM Cloud, including security, compliance, and auditability requirements.

Security and compliance

Solution infrastructure deployed on IBM Cloud following IBM Cloud for Financial Services framework must adhere to the regulations and security controls that govern secure software supply chain practices, software integrity, boundary protection, and outbound traffic control policies, as applicable to container image management:

Container images should only be stored on authorized managed registries

The use of unauthorized registries to store images must be forbidden (except for the tools responsible to internalize images). See SI-7 and CM-6.

Network access to external resources on public network should be restricted to a predefined allow-list

See SI-4 (4), SC-7, and SC-7 (5).

Container images should be checked for integrity and vulnerabilities

See SI-7 and CM-6.

Auditability

The container image management process should provide a clear picture of the image inventory deployed within the IBM Cloud solution infrastructure and the state of their vulnerability exposure.

Mitigating security concerns

Security risks such as access to unauthorized container registries, deployment of container images without proper security checks, or newly discovered vulnerabilities must be mitigated. Specific access controls should be applied to the processes that facilitate mirroring of external registries and making the container images available to the deployed clusters in IBM Cloud.

Secure software development

For container images built by the solution operators, a secure software development environment such as IBM DevSecOps with Continuous Delivery is recommended to ensure end-to-end security of the custom software components.

Decision matrix

Use the following table to help determine the best approach for your workload:

Criteria	Questions to ask	Recommended option

Containerized workload deployment process	Does your workload consist of custom container images with full control over deployment manifests?	Publish your images to private Container Registry namespace, modify image references in deployment manifests to point to Container Registry. You may not need registry mirroring configuration on cluster nodes.
Use of public container images for deployments	Do you use container images from public registries in your own deployment manifests?	Set up image mirroring process for the target images, modify image references in deployment manifests to point to Container Registry. You may not need registry mirroring configuration on cluster nodes.
Use of Red Hat OpenShift on IBM Cloud operators from public marketplaces	Do you need to install operators from public marketplaces?	Set up image mirroring process for the target operators, configure registry mirroring and registry credentials on cluster nodes using DaemonSets.
Access to image registries on private networks	Do you deploy workloads using images from registries on private networks such as enterprise internal registries?	Set up image mirroring process running within a private network that also has a secure connection to IBM Cloud and mirror the required images to a private namespace in Container Registry. Modify image references in deployment manifests to point to Container Registry. You may not need registry mirroring configuration on cluster nodes.

Decision matrix for image management approaches

Next steps

- Start implementing your image management strategy with Container Registry configuration

Configuring IBM Cloud Container Registry

Implementation guidance for control CM-6 - Configuration Settings requires that all container images are stored in authorized managed image registries. IBM Cloud Container Registry with private namespaces should be used for the workloads deployed on your Red Hat OpenShift on IBM Cloud cluster.

Before you begin

Ensure that your cluster VPC infrastructure already implements the boundary protection measures and outbound traffic restrictions:

- Public gateways should be deployed only on VPC subnets where the workloads deployed on cluster worker nodes require outbound access to allow-listed specific external endpoints. Direct access to external registries should not be allowed from cluster worker nodes.
- No routes in the VPC that would allow unrestricted egress to public networks from cluster nodes.
- ACL for cluster subnets do not allow unrestricted connections outside of VPC address space and IBM Cloud services (no CIDR 0.0.0.0/0).

Creating and organizing Container Registry namespaces

When creating new private namespaces in Container Registry , assign them to resource groups based on their security or workload classification. This will help to organize IAM access policies that may be using resource groups for qualification.

- Avoid using Default resource group or resource groups used for purposes not related to the images to be stored in the namespace.
- Consider creating separate namespace(s) for operator images if you are planning to mirror them.

Setting up staging and production namespaces

Create separate Container Registry namespaces using appropriate naming conventions for Staging vs. Production image sets and use separate mirroring processes to update them.

1. Ensure that the images in Staging namespaces are scanned for vulnerabilities and tested for update impact before mirroring them to production namespaces.
2. Set up separate mirroring processes for copying images to Staging namespace from the source and then mirroring from Staging to Production based on validation results using a separate Image Set configuration.

Configuring access control

Creating Service IDs for image pulling

Create IBM Cloud IAM Service IDs with limited access (for example, Reader service access) scoped to specific namespaces to be used in pull secret configuration.

This will ensure that minimally required access is granted for pulling the images by the cluster. The default cluster pull secrets include credentials for a cluster's IBM Cloud IAM Service ID with reader access to your Container Registry namespaces.

Creating Service IDs for image mirroring

Create separate IBM Cloud IAM Service IDs for mirroring process that will be used to push images to your private namespaces in Container Registry.

1. Assign Writer service role, optionally scoped to specific namespaces.
2. These API keys should only be used in the processes that perform the image mirroring, typically outside of the Red Hat OpenShift on IBM Cloud cluster.

Setting up network access

Creating Virtual Private Endpoints

Create Virtual Private Endpoints for controlled access to Container Registry from the compute resources in your VPCs.

Configuring Context Based Restrictions

Create Context Based Restrictions for Container Registry to enforce access over private endpoints and specific source networks.

Using Container Registry with OpenShift

Once you have the Container Registry namespaces and proper access control configured, you can set up your Red Hat OpenShift on IBM Cloud cluster to use Container Registry instead of the internal registry.

For more information, see Using IBM Cloud Container Registry with OpenShift.

For the deployments where you can modify the container image reference, you will need to update the image location in the deployment manifest to point to the image that you pushed to your private Container Registry namespace.

Managing images in Container Registry

The mirroring process will not remove images from the target Container Registry namespace even if they are not in the Image Set configuration for mirroring. Therefore, you will need to manually remove images that are no longer used. You should also remove obsolete image versions with identified vulnerabilities.

Next steps

- Learn how to copy images to Container Registry with oc-mirror CLI
- Configure registry mirroring on Red Hat OpenShift on IBM Cloud nodes

Copying images to Container Registry with oc-mirror CLI

The `oc-mirror` plugin can be used for mirroring images to private Container Registry from external registries. You can run the `oc-mirror` plugin using either a VSI or as a Code Engine job.

Before you begin

Create a new IBM Cloud IAM Service ID that is granted Reader and Writer access to your Container Registry private namespaces that will be used for the mirrored images. In the following examples, `container-image-mirror-service` is used as the Service ID name.

Option 1: Running oc-mirror plugin using a VSI

To use the `oc-mirror` plugin to mirror registry images, you must install the plugin. If you are mirroring image sets in a fully disconnected environment, ensure that you install the `oc-mirror` plugin on the host with internet access and the host in the disconnected environment with access to the mirror registry.

The credentials for the Container Registry access will be stored in a local file.

Installing the oc-mirror plugin

1. Download and install the `oc-mirror` plugin following the instructions in the Red Hat documentation.
2. Make sure you have `ibmcloud` CLI, `oc` CLI, and `jq` tool installed. The command examples are for bash shell.

Setting up authentication

Run the following commands to create a new API key for the IBM Cloud IAM Service ID, pull the Red Hat credentials from the cluster to a local file `image-pull-secret.json`, and add the Container Registry credentials (API key created) to it. You need to log in to IBM Cloud with `ibmcloud` CLI before running the commands. Update the values in the export commands as needed to match your environment.

```
export SERVICE_ID_NAME=container-image-mirror-service
export REGISTRY=private.de.icr.io
export CLUSTER_NAME=workload-cluster

apikey=$(ibmcloud iam service-api-key-create oc-mirror-icr-access-key-vsi \
"$SERVICE_ID_NAME" -d 'Credentials for ICR access from VSI' --output json \
| jq -r '.apikey')

ibmcloud oc cluster config -c $CLUSTER_NAME --admin

oc get secret -n openshift-config pull-secret \
-o template='{{index .data ".dockerconfigjson"}}' | base64 --decode \
| jq --arg registry "$REGISTRY" \
--arg auth "$(echo -n iamapikey:$apikey | base64 -w0)" \
'.auths[$registry].auth = $auth' > image-pull-secret.json
```

If you need to access other registries not included in your cluster's pull secret, edit the `image-pull-secret.json` file to add the authentication values for these registries. For more information, see Adding mirror registry authentication details.

Creating ImageSetConfiguration

Create an `ImageSetConfiguration` YAML file with the list of operators and images to be mirrored. The following example shows a sample `ImageSetConfiguration`. For more details on operator versions and channels to be mirrored, see Defining and configuring your operators for use with oc-mirror. Save the file as `mirror-imageset.yaml`.

```
kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v2alpha1
mirror:
  platform:
  operators:
    - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.20
      packages:
        - name: volsync-product
        - name: cluster-observability-operator
    - catalog: quay.io/operatorhubio/catalog:latest
      packages:
        - name: keycloak-operator
  additionalImages:
  - name: quay.io/keycloak/keycloak:26.5.6
  - name: registry.redhat.io/ubi9-minimal:latest
```

In the Image Set configuration, you will need to list the catalog images for the catalogs you intend to use. For example, the OperatorHub.io catalog is published on `quay.io/operatorhubio/catalog:latest`. Under each catalog entry, you can list the operators (as packages) that you need to mirror. You can configure multiple catalog sources in the single Image Set.

Some operators require additional images to be pulled as they do not provide complete operator bundles. Those images should be listed under `additionalImages`.

Running oc-mirror

Run the `oc-mirror` v2 command to start the mirroring to target registry. Make sure to add `umask 0022` to skip cache issues for index image. Update the values in the export commands as needed to match your environment. Make sure the namespace used as `TARGET_NAMESPACE` has been created as described in Configuring IBM Cloud Container Registry.

```
export REGISTRY=private.de.icr.io
export TARGET_NAMESPACE=ocp-operator-mirror

oc-mirror --config mirror-imageset.yaml --workspace file://oc-mirror-v2 \
docker://$REGISTRY/$TARGET_NAMESPACE --v2 --authfile image-pull-secret.json
```

The `oc-mirror` plugin automatically creates the index images with the subset of operators in the target registry and CatalogSource in the `oc-mirror-v2` directory. The CatalogSource content will be used later on the cluster to allow operator installation.

Option 2: Running oc-mirror plugin as a Code Engine job

Creating a Code Engine project

Create a new serverless project for Code Engine using the following command:

```
$ ibmcloud ce project create --name <project-name>
```

Creating a registry secret

Create a new registry secret by generating a new API key for the IBM Cloud IAM Service ID configured in the prerequisites section:

```
export SERVICE_ID_NAME=container-image-mirror-service
export REGISTRY=private.de.icr.io

apikey=$(ibmcloud iam service-api-key-create oc-mirror-icr-access-key-ce \
"$SERVICE_ID_NAME" -d 'Credentials for ICR access from CE' --output json \
| jq -r '.apikey')

ibmcloud ce registry create --name icr-secret \
--server $REGISTRY --username iamapikey --password-from-file <(echo -n "$apikey")
```

Setting up the build configuration

Set up the build configuration in the Code Engine project by using the following command. The Dockerfile used for the image build is available in the GitHub repository.

It is recommended to have this Dockerfile in a private registry and update the references in the following command based on the private repository path. SSH key also needs to be added as an additional parameter to the command while using the private repository.

```
ibmcloud ce build create --name oc-mirror-cli-image-build \
--build-type git \
--source https://github.com/IBM/fs-solution-guides-resources.git \
--context-dir container-image-management/oc-mirror-ce \
--dockerfile ce-image.dockerfile \
--image $REGISTRY/oc-mirror-tools/oc-mirror-tool:latest
```

Building the custom image

Run the following build command to create the custom image and push it to Container Registry:

```
ibmcloud ce buildrun submit --build oc-mirror-cli-image-build
```

Creating authentication secret

Run the following commands to create a new API key for the IBM Cloud IAM Service ID, pull the Red Hat credentials from the cluster to a local file `image-pull-secret.json`, and add the Container Registry credentials (API key created) to it. Update the values in the export commands as needed to match your environment.

```
export SERVICE_ID_NAME=container-image-mirror-service
export REGISTRY=private.de.icr.io
export CLUSTER_NAME=workload-cluster

apikey=$(ibmcloud iam service-api-key-create oc-mirror-icr-pull-secret-key-ce \
"$SERVICE_ID_NAME" -d 'Credentials for ICR access from CE' --output json \
| jq -r '.apikey')

ibmcloud oc cluster config -c $CLUSTER_NAME --admin

icr_auth=$(oc get secret -n openshift-config pull-secret \
-o template='{{index .data ".dockerconfigjson"}}' | base64 --decode \
| jq --arg registry "$REGISTRY" \
--arg auth "$(echo -n iamapikey:$apikey | base64 -w0)" \
'.auths[$registry].auth = $auth')

ibmcloud ce secret create --name oc-mirror-auth \
--from-file image-pull-secret.json=<(echo -n "$icr_auth")
```

Creating ConfigMap for ImageSetConfiguration

Create the `ImageSetConfiguration` YAML file containing the images to be mirrored (see step 3 in Option 1) and create a new ConfigMap from the `ImageSetConfiguration` YAML using the following command. For more details on `ImageSetConfiguration` YAML structure and parameters, see Defining and configuring your operators for use with oc-mirror.

```
ibmcloud ce configmap create \
  --name oc-mirror-config \
  --from-file=mirror-imageset.yaml
```

Creating the Code Engine job

Create a new job and pass the secret, ConfigMap, and arguments using the following command:

```
export REGISTRY=private.de.icr.io
export TARGET_NAMESPACE=ocp-operator-mirror

ibmcloud ce job create \
  --name oc-mirror-job \
  --image $REGISTRY/oc-mirror-tools/oc-mirror-tool:latest \
  --registry-secret icr-secret \
  --cpu 2 \
  --memory 8G \
  --ephemeral-storage 4G \
  --mount-secret /auth=oc-mirror-auth \
  --mount-configmap /workspace=oc-mirror-config \
  --command oc-mirror \
  --argument --config=/workspace/imageset.yaml \
  --argument --workspace=file:///oc-mirror-v2 \
  --argument --authfile=/auth/config.json \
  --argument docker://$REGISTRY/$TARGET_NAMESPACE \
  --argument --v2
```

Setting up persistent storage

1. Create a new Cloud Object Storage bucket and HMAC credentials to be used for Code Engine for enabling the persistent storage by following the documentation. The bucket contains the results of the mirroring process such as generated CatalogSource files as well as some temporary files created by `oc-mirror` during registry replication.
2. Take a note of the Access Key ID and Secret Access Key from the credentials created in the previous step and create a new HMAC secret in the Code Engine project.
3. Create a new persistent data store in the Persistent data stores section on the Code Engine project page. Use a descriptive name like `oc-mirror-cos-workdir`, select COS instance and bucket used in step 1 and an HMAC secret created in step 2.
4. In the Volume Mounts section in Code Engine job created, add the persistent data stores and the mount path in the COS bucket (`/oc-mirror-v2`). This is the location where `oc-mirror` plugin stores all the cache files and generated YAML files during mirroring.

Running the job

Run the job created using the following command:

```
ibmcloud ce jobrun submit --job oc-mirror-job
```

The `oc-mirror` plugin automatically creates the index images with the subset of operators in the target registry and CatalogSource to be deployed to the cluster in the `oc-mirror-v2` directory inside the COS bucket.

Next steps

- Configure registry mirroring on Red Hat OpenShift on IBM Cloud nodes
- Set up operator mirroring and catalog sources

Configuring registry mirroring on Red Hat OpenShift on IBM Cloud nodes

The CRI-O container engine on Red Hat OpenShift on IBM Cloud nodes uses configuration files to resolve the requests for pulling an image from a registry. The configuration can be updated to include an alternative location.

Understanding registry mirroring

Consider the following example for the registry configuration on a node:

```
[[registry]]
prefix = "quay.io/example"
insecure = false
blocked = false
location = "private.icr.io/acme-workloads"

[[registry.mirror]]
location = "private.de.icr.io/acme-workloads"
```

When a pull of `quay.io/example/image:latest` is requested by a deployment, the node will request the image from `private.de.icr.io/acme-workloads/image:latest` instead, thus avoiding the need for a direct connection to an external public registry (`quay.io`).

The access to Container Registry requires valid IBM Cloud API key which can be provided at a namespace or deployment level as a pull secret in case of regular workload deployment manifests. However, for operator deployments and catalog mirroring, the pull secret must be at the cluster level and needs to be distributed to nodes via a DaemonSet.

Once the registry mirroring and pull secrets are configured on the nodes, you will not need to modify your deployment manifests for image references from external registries. Instead, the images with references to those registries will be pulled from specified Container Registry locations.

Before you begin

IBM Cloud does not currently support updating the `ImageContentSourcePolicy`, `ImageDigestMirrorSet`, and `ImageTagMirrorSet`. Therefore, DaemonSets can be used to emulate the above-mentioned policies on the nodes.

Once the images are mirrored to the target registry (Container Registry) using `oc-mirror` plugin, two DaemonSets need to be created on the cluster.

Creating a DaemonSet for registry configuration update

Creating the mapping file

Create a new `mapping.txt` file with source and mirror references for the registries and images that you copied to Container Registry, using the following pattern:

```
registry.redhat.io=private.de.icr.io/mirror-registry
quay.io=private.de.icr.io/mirror-registry
```

Creating the ConfigMap

Create a new ConfigMap using the `mapping.txt` in the `kube-system` namespace using the following command:

```
oc create configmap registry-mapping --from-file mapping.txt -n kube-system
```

Deploying the DaemonSet

Deploy a new DaemonSet for updating the mirror registry configuration to `registries.conf.d` directory in the node. Use the sample YAML in the GitHub repository for creating the DaemonSet.

Creating a DaemonSet for updating node level Container Registry pull secret

Creating the pull secret

Create a new secret using the API key credentials created earlier using the following command:

```
oc create secret docker-registry docker-auth-secret \
--docker-server=private.de.icr.io \
--docker-username=iamapikey \
--docker-password=<password> \
--namespace kube-system
```

Deploying the DaemonSet

Create a new DaemonSet for updating the Container Registry pull secret in the node. The sample YAML is available in the GitHub repository that can be used for deployment.

Updating DaemonSets

DaemonSets need to be restarted by using the following command whenever there is an update to secret or ConfigMap created in the previous steps:

```
oc rollout restart ds <daemonset-name> -n <namespace>
```

Next steps

- Configure operator mirroring and catalog sources
- Set up image acceptance policies

Configuring operator mirroring and catalog sources

Since operators are deployed via images, the solution to deploy operators originating from external registries involves the same approach as image mirroring with a few additions.

Understanding operator mirroring

Without a direct access to external catalogs, catalog index images are re-created using `oc-mirror` command line tool. Configuring CatalogSource resources with the repackaged index images will make the selected operators visible in the Red Hat OpenShift on IBM Cloud marketplace view and enable their installation.

To facilitate the proper operator images deployment on an isolated cluster, you need to have the registry configuration and node pull secret credentials update as described in Configuring registry mirroring on Red Hat OpenShift on IBM Cloud nodes .

Creating CatalogSource resources

The `oc-mirror` plugin automatically creates the YAML file for deploying the CatalogSource in the directory `/oc-mirror-v2/working-dir/cluster-resources`. The following example shows a sample CatalogSource YAML created by `oc-mirror`:

```
apiVersion: operators.coreos.com/v1alpha1
kind: CatalogSource
metadata:
  annotations:
    createdAt: Friday, 30-Jan-26 06:15:05 UTC
    createdBy: oc-mirror v2
    oc-mirror_version: 4.20.0-202512150315.p2.gf4775a2.assembly.stream.el9-f4775a2
  name: cs-redhat-operator-index-v4-20
  namespace: openshift-marketplace
spec:
  image: private.de.icr.io/mirror-registry/redhat/redhat-operator-index:v4.20
  sourceType: grpc
status: {}
```

Applying the CatalogSource

This YAML can be applied to the cluster using the following command, which would list the mirrored operators in the OperatorHub:

```
oc apply -f <file-name>
```

Operators can then be installed following the normal operator installation process.

Verifying operator availability

After applying the CatalogSource, you can verify that the operators are available in the Red Hat OpenShift on IBM Cloud console:

1. Navigate to **Operators** > **OperatorHub** in the Red Hat OpenShift on IBM Cloud console.
2. Search for the operators you mirrored.
3. Verify that the operators appear in the list and can be installed.

Next steps

- Configure image acceptance policies
- Set up vulnerability scanning

Configuring image acceptance policies

Portieris is a Kubernetes admission controller that verifies container images before deploying them to the cluster. You can create image security policies for

each namespace or at the cluster level and enforce different rules for different images or namespaces.

If an image doesn't meet the defined policy requirements, the resource that contains the pod is not deployed to the cluster.

Installing Portieris

Portieris can be installed by toggling the Image security enforcement option to 'Enabled' from the cluster overview page in IBM Cloud Console or using CLI by following the documentation.

Understanding policy types

For system namespaces (except default namespace) which are automatically created during cluster creation, Image acceptance policies are automatically defined and controlled by Portieris and cannot be overridden by the users.

For custom or new namespaces created by users in the cluster, Portieris Image Acceptance Policies can be enforced to ensure that only whitelisted images are deployed into the pods of the cluster. Portieris defines two custom resource types for policy:

Cluster image policy resources

> `ClusterImagePolicy` resources are configured at the cluster level and are not specific to any namespace. They take effect whenever an image policy resource, `ImagePolicy`, is not defined in the namespace where the workload is deployed. If a matching policy is not found for an image, deployment is denied.

Image policy resources

> `ImagePolicy` resources are configured in the target namespace and define Portieris' behavior in that namespace. If image policy resources exist in a namespace, the policies from those image policy resources are used exclusively. If a match doesn't exist for a workload image in `ImagePolicy`, cluster image policy resources, `ClusterImagePolicy`, are not examined.

Creating a default deny-all cluster policy

By default, `ClusterImagePolicy` would allow all the images to new or custom namespaces, and it should be changed to deny all by default so that no random images can be deployed to the cluster.

The following example shows a default policy which denies all the deployment to the cluster. This would be applicable to all the new or custom namespaces:

```
apiVersion: portieris.cloud.ibm.com/v1
kind: ClusterImagePolicy
metadata:
  name: block-all-images
spec:
   repositories: []
```

If debug pods need to be run, the debug pod image also has to be whitelisted in the `ClusterImagePolicy`. Debug pod images can be referenced as `quay.io/openshift-release-dev/ocp-v4.0-art-dev*`.

Creating namespace-specific image policies

Images are matched against the set of policies defined. If a policy doesn't match the workload image, the deployment of a container is denied. This ensures that only whitelisted images are allowed and denies all other deployment to cluster. The policy should match the original image location, even if the image is pulled from a mirror.

The following example allows any image from a specific repository only and denies deployment from any other repository:

```
apiVersion: portieris.cloud.ibm.com/v1
kind: ImagePolicy
metadata:
  name: test-policy-portieris
  namespace: test-namespace
spec:
   repositories:
    - name: "private.de.icr.io/mirror-registry/*"
    - name: "quay.io/keycloak/keycloak*"
    - name: "registry.redhat.io/ubi-minimal:latest"
      policy:
```

Verifying image signatures

To verify image signatures using GPG public keys, you will need to create a secret containing the GPG public key and reference it in the `ImagePolicy` specification:

```
apiVersion: portieris.cloud.ibm.com/v1
kind: ImagePolicy
metadata:
  name: test-policy-portieris
  namespace: test-namespace
spec:
   repositories:
    - name: 'icr.io/appcafe/open-liberty:kernel-slim-java17-openj9-ubi-amd64'
      policy:
        simple:
          requirements:
            - keySecret: open-liberty-signed-image
              type: signedBy
```

In this example, an OpenLiberty image is used that is signed by the key provided on the OpenLiberty project page. The opaque secret `open-liberty-signed-image` needs to be created in the same namespace and have a key "key" with data containing the GPG public key as a text.

Example: Allowing Sysdig agent for Security and Compliance Center Workload Protection

The following example shows how to allow installation of Sysdig agent for Security and Compliance Center Workload Protection:

```
apiVersion: portieris.cloud.ibm.com/v1
kind: ImagePolicy
metadata:
  name: allow-sysdig-policy
  namespace: ibm-observe
spec:
  repositories:
    - name: 'icr.io/ext/sysdig/*'
```

Next steps

- Set up vulnerability scanning

Scanning container images for vulnerabilities in IBM Cloud

Container images should be periodically scanned for vulnerabilities to ensure the security of your workloads. This topic describes the available scanning options for images in Container Registry and deployed on Red Hat OpenShift on IBM Cloud clusters.

SCC WP registry scanning

Security and Compliance Center Workload Protection (SCC WP) requires a registry scanning process using a pre-built Sysdig image to be set up with access to Container Registry. The scan results are reported back to SCC WP.

For details on how to set up a scanning job on IBM Cloud Code Engine, see Scanning vulnerabilities with the IBM Cloud Security and Compliance Center Workload Protection.

SCC WP agent for scanning of images deployed on a cluster

By deploying a Sysdig agent for SCC WP on an Red Hat OpenShift on IBM Cloud cluster, you enable scanning of all the images of the deployed containers for vulnerabilities.

The agent deployment can be triggered from the cluster overview in IBM Cloud console by enabling Workload Protection integration. It is recommended to have a fully configured SCC WP and Cloud Monitoring instances configured per IBM Cloud for Financial Services reference architecture before enabling the integration for a cluster.

Make sure you add corresponding image acceptance policies for Sysdig observability components. For an example, see Allowing Sysdig agent for Security and Compliance Center Workload Protection.

Container Registry scanning

Container images uploaded to Container Registry are automatically scanned for vulnerabilities. However, the scan results are not reported to SCC WP and would have to be evaluated separately. It is recommended to set up a scanning process using SCC WP image scanner as described in SCC WP registry

scanning for consistent vulnerability reporting.

Next steps

- Check the references for additional resources

Container image management references

Explore additional resources for container image management, IBM Cloud container registry, and related topics, including documentation and external links.

IBM Cloud documentation

IBM Cloud Container Registry documentation

IBM Cloud Container Registry overview

Using IBM Cloud Container Registry with OpenShift

Using IBM Cloud Container Registry with OpenShift

Signing images in Container Registry

Signing images for trusted content

Red Hat Red Hat OpenShift on IBM Cloud documentation

Disconnected installation mirroring

Red Hat OpenShift Container Platform - Disconnected installation mirroring

Installing oc-mirror plugin

Installing the oc-mirror plugin

Defining and configuring operators for oc-mirror

Defining and configuring your operators for use with oc-mirror

IBM Cloud for Financial Services

IBM Cloud for Financial Services framework controls

DevSecOps with Continuous Delivery

Additional resources

GitHub repository for solution guides

IBM Financial Services solution guides resources

Security and Compliance Center Workload Protection

Scanning vulnerabilities with IBM Cloud Security and Compliance Center Workload Protection

OpenLiberty image signature verification

Verify signatures for container images in Open Liberty

Storage

Storage for VPC reference architecture

Learn about the Financial Services Validated storage options that are used within the VPC reference architecture.

IBM® Cloud Block Storage for Virtual Private Cloud

Block Storage for VPC provides hypervisor-mounted, high-performance data storage for your virtual server instances that you can provision within a VPC. The VPC infrastructure provides rapid scaling across zones and extra performance and security.

Block Storage for VPC is used for both primary boot volumes and secondary data volumes. Boot volumes are automatically created and attached during instance provisioning. Data volumes can be created and attached during instance provisioning as well, or as stand-alone volumes that you can later attach to an instance. To protect your data, you should use KYOK encryption with Hyper Protect Crypto Services.

To learn more and to start creating Block Storage for VPC volumes, see the following resources:

- About Block Storage for VPC
- Create and attach a block storage volume when you create a new instance
- Create a stand-alone block storage volume
- Creating block storage volumes with customer-managed encryption

IBM Cloud® Object Storage

Object Storage stores encrypted and dispersed data across multiple geographic locations. Object Storage is available with three types of resiliency: Cross Region, Regional, and Single Data Center. Cross Region provides higher durability and availability than using a single region at the cost of slightly higher latency. Regional service reverses those tradeoffs, and distributes objects across multiple availability zones within a single region. If a given region or availability zone is unavailable, the object store continues to function without impediment. Single Data Center distributes objects across multiple machines within the same physical location.

Users of Object Storage refer to their binary data, such as files, images, media, archives, or even entire databases as objects. Objects are stored in a bucket, the container for their unstructured data. Buckets contain both inherent and user-defined metadata. Finally, objects are defined by a globally unique combination of the bucket name and the object key, or name.

> ⚠ **Important:** All Object Storage buckets must be encrypted with KYOK by using keys that are managed by Hyper Protect Crypto Services. For more information, see Encryption at rest. In addition, a geographically separate region should be used as an alternative storage site. This means you should use cross region resiliency for all of your Object Storage buckets.

To start working with Object Storage, see the following instructions:

- About Object Storage
- Getting started with Object Storage
- Connecting to Object Storage from VPC
- Encrypting your data

Related controls in IBM Cloud Framework for Financial Services

More details on some of the major IBM Cloud Framework for Financial Services controls related to storage can be found in the following articles:

- Encryption at rest
- Backup and recovery

Next steps

- Audit logging of IBM Cloud events

Storage for Satellite reference architecture

You need to consider storage options both in IBM Cloud and in the on-premises Satellite location.

Storage in IBM Cloud

All Satellite control plane data is backed up to an IBM Cloud Object Storage service instance within IBM Cloud. This is done so that your location can be

restored after a disaster. When you create a location, you also provide a Object Storage service instance that you control for backup of the location control plane worker nodes. Control plane master data is backed up by IBM and stored in an IBM-owned Object Storage instance. Satellite cluster master data is backed up to the Object Storage instance that you own.

Aside from backing up data from the Satellite location, Object Storage is also used to store auditable IBM Cloud events. For more information, see Audit logging of IBM Cloud events.

About Object Storage

Object Storage stores encrypted and dispersed data across multiple geographic locations. Object Storage is available with three types of resiliency: Cross Region, Regional, and Single Data Center. Cross Region provides higher durability and availability than using a single region at the cost of slightly higher latency. Regional service reverses those tradeoffs, and distributes objects across multiple availability zones within a single region. If a given region or availability zone is unavailable, the object store continues to function without impediment. Single Data Center distributes objects across multiple machines within the same physical location.

Users of Object Storage refer to their binary data, such as files, images, media, archives, or even entire databases as objects. Objects are stored in a bucket, the container for their unstructured data. Buckets contain both inherent and user-defined metadata. Finally, objects are defined by a globally unique combination of the bucket name and the object key, or name.

⚠ **Important:** All Object Storage buckets must be encrypted with KYOK by using keys that are managed by Hyper Protect Crypto Services. For more information, see Encryption at rest. In addition, a geographically separate region should be used as an alternative storage site. This means you should use cross region resiliency for all of your Object Storage buckets.

To start working with Object Storage, see the following instructions:

- About Object Storage
- Getting started with Object Storage
- Connecting to Object Storage from VPC
- Encrypting your data

Storage in Satellite location

Within the Satellite location, Satellite storage uses Satellite Config to provide a convenient way to install various storage drivers in Red Hat OpenShift on IBM Cloud clusters, by using storage templates. The storage templates are provided and tested by the vendors. After you install Satellite storage, your cluster users can use Kubernetes persistent volume claims (PVCs) to order and save their application data in persistent storage. For more information, see Understanding Satellite storage.

⚠ **Important:** You are responsible for providing the underlying virtual and physical storage that will be accessed by Satellite storage in your on-premises Satellite location. You should ensure your underlying physical storage is encrypted using keys managed by an on-premises FIPS 140-2 level 2 or higher hardware security module (HSM).

Related controls in IBM Cloud Framework for Financial Services

More details on some of the major IBM Cloud Framework for Financial Services controls related to storage can be found in the following articles:

- Encryption at rest
- Backup and recovery

Next steps

- Audit logging of IBM Cloud events

Logging and monitoring

Audit logging of IBM Cloud events

When you work in a cloud environment, such as the IBM Cloud, you must plan for auditing and monitoring of workloads and data. This should be guided by your internal policies together with industry and location-based compliance requirements. It’s critical to capture audit events generated by the cloud services and hosted workloads running on the cloud platform. With effective audit logging you can investigate abnormal activity and demonstrate compliance with legal and regulatory requirements. By actively monitoring audit events, you can identify and investigate potential security problems and take corrective action.

A core part of the auditing and logging capabilities for IBM Cloud is the forwarding of audit records for each IBM Cloud service through Activity Tracker Event Routing. Through this service, there is the forwarding of all audit events that record the activity of the IBM Cloud APIs and all consumers of the API, such as, the IBM Cloud CLI and console.

Activity Tracker Event Routing receives all audit events from cloud services and they're forwarded to one of three different targets:

- IBM Cloud Object Storage
- IBM Cloud Event Streams
- IBM Cloud Activity Tracker hosted event search

The architecture for IBM Cloud Activity Tracker Event Routing supports high availability and is hosted in selected regions.

How you process the audit event logs will depend on the data categorization of the audit events and risk appetite set by your organization. Audit event logs are designed to contain information that many organizations would not categorize as extremely sensitive. However, they do contain IP addresses and the names of cloud resources that some organizations might deem as sensitive requiring enhanced assurance through the use of IBM Cloud for Financial Services Validated services.

Activity Tracker Event Routing, IBM Cloud Object Storage and IBM Cloud Event Streams (Enterprise Plan) are all IBM Cloud for Financial Services Validated. It's recommended that you use these services for capturing and forwarding IBM Cloud platform audit events with a high data categorization.

If you organization deems the audit events aren't the highest data classification, they may decide the additional assurance of a service that's IBM Cloud for Financial Services Validated isn't required. In this case, it may be that IBM Cloud Activity Tracker hosted event search offering is an option for your organization.

Understand the handling requirements for audit logs in your organization to decide on the appropriate solution.

Activity Tracker Event Routing

Use Activity Tracker Event Routing to forward auditable events from IBM Cloud services to either IBM Cloud Object Storage or Event Streams (Enterprise Plan).

> **Tip:** Only Event Streams using the Enterprise Plan is IBM Cloud for Financial Services Validated. For more information on Event Streams compliance, see Understanding compliance for Event Streams.

You will find further information on Activity Tracker Event Routing in the cloud documentation:

- A description of Activity Tracker Event Routing
- Configuring for Cloud Object Storage target
- Configuring for Event Streams target
- The Cloud services generating audit events

You must:

- Ensure encryption of object storage with KYOK or BYOK using keys managed either by Hyper Protect Crypto Services or Key Protect, based on the audit events data categorization and risk appetite of your organization. For information on encrypting Cloud Object Storage, see Server-Side Encryption with Hyper Protect Crypto Services.
- Configure storage capacity for audit log storage to support retention requirements with retention of at least 90 days of online records and one years’ worth of records offline to support archival. For managing retention of data, see Deleting stale data with expiration rules.
- Monitor to ensure that any limits, such as quota, aren't reached for Object Storage and are increased if necessary. For setting quotas on storage, see Setting a quota on a bucket.
- Ensure audit events are available from collection from Cloud Object Storage for log analysis or forwarded from Event Streams to your security information and event management (SIEM) solution for threat monitoring.

Tip: Online retention refers to security audit logs that are available for analysis within one hour or less, and offline retention refers to security audit logs that are available for analysis within two business days or less.

Activity Tracker hosted event search

Use IBM Cloud Activity Tracker hosted event search to capture a record of your IBM Cloud activities and monitor the activity of your IBM Cloud account. You can use this service to investigate abnormal activity and critical actions, and comply with regulatory audit requirements. The events that are collected comply with the Cloud Auditing Data Federation (CADF) standard and stored in IBM Cloud Object Storage.

As the service doesn't have Financial Services Validation, we don't recommend it for sensitive or high category data than you may decide to use this capability if the data isn't at your highest level of categorization for your organization, as it offers additional functionality including the ability to display, search and alert on audit events without the need to hosting a log management solution. It may also be a suitable solution for application testing or a proof of concept as the events can be redirected using Activity Tracker Event Routing when the application uses more sensitive data.

You will find further information on Activity Tracker hosted event search in the cloud documentation:

- A description of Activity Tracker hosted event search
- Provisioning an instance of Activity Tracker
- The Cloud services generating audit events

You must:

- Ensure the storage is encrypted with KYOK or BYOK using keys managed either by Hyper Protect Crypto Services or Key Protect, based on the audit events data categorization and risk appetite of your organization. For Activity Tracker hosted event search, see Encrypt data with your own key for more information.
- Configure storage capacity for audit log storage to support retention requirements with retention of at least 90 days of online records and one years' worth of records offline to support archival. For managing retention of data, see Deleting stale data with expiration rules.
- Monitor to ensure that any limits, such as quota, aren't reached for Object Storage and are increased if necessary. For setting quotas on storage, see Setting a quota on a bucket. For managing large numbers of unexpected audit events, see Managing volume spike protection and costs.
- Ensure audit events are configured to forward to Event Streams and onto to your security information and event management (SIEM) solution for threat monitoring.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-2 Account Management AC-2 (1) Account Management | Automated System Account Management AC-2 (4) Account Management | Automated Audit Actions AC-2 (7) Account Management | Privileged User Accounts
Audit and Accountability (AU)	AU-3 Content of Audit Records AU-4 Audit Log Storage Capacity AU-5 Response to Audit Logging Process Failures AU-6 Audit Record Review, Analysis, and Reporting AU-6 (1) Audit Record Review, Analysis, and Reporting | Automated Process Integration AU-7 Audit Record Reduction and Report Generation AU-10 Non-repudiation AU-11 Audit Record Retention

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management AC-2 (1) Account Management | Automated System Account Management AC-2 (4) Account Management | Automated Audit Actions AC-2 (7) Account Management | Privileged User Accounts

Audit and Accountability (AU)	AU-3 Content of Audit Records AU-4 Audit Log Storage Capacity AU-5 Response to Audit Processing Failures AU-6 Audit Record Review, Analysis. and Reporting AU-6 (1) Audit Record Review, Analysis. and Reporting | Automated Process Integration AU-7 Audit Record Reduction and Report Generation AU-10 Non-repudiation AU-11 Audit Record Retention

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Audit logging of application provider events and SIEM

Audit logging of application provider events and SIEM

In addition, to audit logging of IBM Cloud events, you also need to consider auditing of events your applications might generate. Auditable events from all sources need to be consolidated into a security information and event management (SIEM) system.

Auditable events

You should audit the following events that occur within the system (in all environments, such as dev, test, prod):

- Access, downloading, revisions to Confidential information
- Access to or update to any financial transaction data
- Any changes to customer accounts
- Login or logoff event, success or failure
- Changes to a system that impact business logic or work flow, examples include but are not limited to, changing call routing, interest rates that are offered for a class of customer, a website flow a customer uses
- Any changes to customer systems that include but not limited to: software install or uninstall, configuration changes, directory or file moves, adds, or deletes, configuration changes, port turn ups, MAC's, routing changes, and so on.
- Each command action taken, both successful and unsuccessful, by a privileged account that includes but not limited to:
    - All changes, additions, or deletions to any account such as lock, unlock, password reset, permission changes,
    - Modification to system time / time-synchronization configuration,
    - Log files: access to logs files; initialization, stopping or pausing of logging, log modification, changes to access permission on log files,
    - Files and system level objects creation and deletion,
    - Attempts to perform unauthorized functions or access data that the user is not authorized to access,
    - Modification of security rules,
    - Application and related systems start-ups and shut-downs
- For privileged accounts:
    - Attempts to execute privilege elevation
    - System errors relevant to security events, including but not limited to SQL errors that indicate an SQL injection, fuzzing, multiple failed logins, failed configuration change, failed/disabled anti-virus software, service failures
    - Web access events in extended log format

The events that are collected should comply with the Cloud Auditing Data Federation (CADF) standard. And, audit events must contain at a minimum:

- Type of event
- When the event occurred
- Where the event occurred
- Source of the event
- Outcome (success or failure) of the event
- Identity of any user/subject/device associated with the event
- Port and protocol, where available
- Session duration, identification, and modification where available
- Authentication method

Storing events in a SIEM

A SIEM system provides the ability to gather security data from information system components and presents that data as actionable information through a single interface. Events of the type in the previous section should be stored in a SIEM. You need to install your own software solution within your VPC for SIEM.

Related IBM Cloud for Financial Services controls

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Access Control (AC)	AC-2 Account Management AC-2 (1) Account Management | Automated System Account Management AC-2 (4) Account Management | Automated Audit Actions AC-2 (7) Account Management | Privileged User Accounts
Audit and Accountability (AU)	AU-3 Content of Audit Records AU-4 Audit Log Storage Capacity AU-5 Response to Audit Logging Process Failures AU-6 Audit Record Review, Analysis, and Reporting AU-6 (1) Audit Record Review, Analysis, and Reporting | Automated Process Integration AU-7 Audit Record Reduction and Report Generation AU-10 Non-repudiation AU-11 Audit Record Retention

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Access Control (AC)	AC-2 Account Management AC-2 (1) Account Management | Automated System Account Management AC-2 (4) Account Management | Automated Audit Actions AC-2 (7) Account Management | Privileged User Accounts
Audit and Accountability (AU)	AU-3 Content of Audit Records AU-4 Audit Log Storage Capacity AU-5 Response to Audit Processing Failures AU-6 Audit Record Review, Analysis. and Reporting AU-6 (1) Audit Record Review, Analysis. and Reporting | Automated Process Integration AU-7 Audit Record Reduction and Report Generation AU-10 Non-repudiation AU-11 Audit Record Retention

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Operational logging

Operational logging

Operational logs are an important complement to audit logs for ensuring your application runs smoothly. Proper operational logging can help you determine whether you need to fail over to an alternative storage or processing site. In addition, operational logging can help you determine whether operations have returned to normal after a system disruption.

Audit logs contain auditable events, while operational logs contain everything else that might be logged. For example, an operational log might contain entries that reflect what's happening as a program executes, such as functions getting called or stack traces getting thrown. While this kind of log data is key to keeping a system running smoothly, it's typically not sent to a SIEM.

Operational logs can be split into two categories:

- Application - Log data generated by software components that you deploy and manage within your deployment. This might be from your own code, or other software components like databases or message queues.
- Platform - Log data from IBM Cloud service instances that is not contained in the audit the data sent to Activity Tracker Event Routing.

Application logging

You need to install your own software solution for capturing application log data within your VPC. There are various ways an operational logging solution can be implemented. See Setting up an operational logging solution for one example that uses the Red Hat OpenShift on IBM Cloud Container Platform Elasticsearch, Fluentd, and Kibana EFK stack.

Platform logging

VPC network traffic

If you are using virtual server instances, you are required to use Flow Logs for VPC to enable the collection, storage, and presentation of information about the Internet Protocol (IP) traffic going to and from network interfaces on virtual server instances within your VPC.

Flow logs can help with a number of tasks, including:

- Troubleshooting why specific traffic isn't reaching an instance, which helps to diagnose restrictive security group rules
- Recording the metadata network traffic that is reaching your instance
- Determining source and destination traffic from the network interfaces
- Adhering to compliance regulations
- Assisting with root cause analysis

See Creating a flow log collector for more details. You must configure flow logs to go to a properly secured IBM Cloud Object Storage bucket, encrypted with KYOK.

Red Hat OpenShift on IBM Cloud and virtual server instances

Platform logs generated from Red Hat OpenShift on IBM Cloud and virtual server instances can be collected by using your own software solution or by using the previously mentioned solution for Setting up an operational logging solution.

Other IBM Cloud services

⚠ **Important:** There is not a Financial Services Validated solution for capturing operational platform logs from other IBM Cloud services. These are usually collected by sending them to IBM Cloud Log Analysis. However, if seeking the Financial Services Validated designation, you should *not* send platform logs to IBM Cloud Log Analysis.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Contingency Planning (CP)	CP-2 (3) Contingency Plan | Resume Mission and Business Functions CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-10 System Recovery and Reconstitution
System and Information Integrity (SI)	SI-11 Error Handling

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Contingency Planning (CP)	CP-2 (3) Contingency Plan | Resume Essential Missions / Business Functions CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-10 Information System Recovery and Reconstitution
System and Information Integrity (SI)	SI-11 Error Handling

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Compliance monitoring

IBM Cloud Logs best practices for Financial Services

Use these best practices to help you use IBM® Cloud Logs features while maintaining Financial Services compliance and reducing costs.

This document is intended for Financial Services compliant ISVs. For general information about using IBM Cloud Logs, see Getting started with IBM Cloud Logs.

Understanding data retention

When you provision IBM Cloud Logs, you specify a retention period that defines how long log data stays in Priority Insights before being moved to archive in IBM Cloud Object Storage. This retention period applies only to Priority Insights, not to all log data.

IBM Cloud Logs Priority Insights retention period

You define the overall retention period for your logs by using Object Storage, not IBM Cloud Logs.

Recommended retention period

For most use cases, retain your logs in Priority Insights for **7 days**. This approach provides fast access and searching of log data for the following scenarios:

- Installing and configuring your infrastructure
- Triaging a Severity 1 or 2 incident
- Developing new features and functions where developers can tail logs live

After 7 days in Priority Insights, logs are automatically moved to Object Storage. Configure the full retention period for logs on the Object Storage bucket when you provision IBM Cloud Logs.

Configuring IBM Cloud Logs storage

After you provision IBM Cloud Logs, create two Object Storage buckets:

Data storage bucket

: Stores the raw log data

Metrics storage bucket

: Stores metrics that are generated from your logs and events

Creating storage buckets

To create the two required buckets:

1. Create a new instance of Object Storage to hold your buckets. For more information, see Provisioning an instance of Object Storage.
2. Open your IBM Cloud Resources page.
3. Expand **Storage** and select your newly created Object Storage instance.
4. Click **Create bucket**.
5. In the **Create a Custom Bucket** section, click **Create**.
6. Enter a meaningful name for your bucket.
7. Verify the correct region is selected in Location.
8. Click **Add+** in the **Archive Rule** section under Advanced Configurations.
9. Set **Time to Archive** to 3 years.
10. Click **Save** at the top of the Archive Rule form.
11. Enable **Key management** under Service ingregrations section, configure the encryption key. For more details on configuring the encryption keys, see Data encryption at rest.

12. Click *_Create bucket_.
13. Repeat these steps to create the second bucket.

Connecting buckets to IBM Cloud Logs

After you create the two buckets, assign them to IBM Cloud Logs:

1. In the IBM Cloud console, click the navigation menu and select **Observability** > **Logging**.
2. Click your IBM Cloud Logs instance.
3. Click **Storage** in the navigation menu.
4. In the **Data Bucket** section, click **Connect**.
5. Select your Object Storage instance and the bucket that you created for data storage.
6. Click **Connect**.
7. In the **Metrics bucket** section, click **Connect**.
8. Select your Object Storage instance and the bucket that you created for metrics storage.
9. Click **Connect**.

Understanding data tiers

IBM Cloud Logs offers three tiers of features: Priority Insights, Analyze and Alert, and Storage and Search. For more information about the features that are supported in each tier, see Data pipelines.

Priority Insights

: Hot storage where logs are indexed immediately and can be searched at high speed. Priority Insights is the most expensive storage tier.

Analyze and Alert

: Medium-priority storage for logs that need alerting and analysis capabilities but not immediate indexing.

Storage and Search

: Archive storage in Object Storage where logs can still be searched, but with slower performance than Priority Insights.

To help lower the cost of log storage, IBM Cloud Logs automatically moves your Priority Insights logs to Object Storage. You can also define TCO policies to control which logs go to Priority Insights and which logs are sent to other pipelines. For more information, see Using the TCO Optimizer.

Data pipeline flow

Each log that you send to IBM Cloud Logs goes through a data pipeline, which defines how that log is:

- Collected
- Processed and enriched
- Stored
- Made available for querying, metrics, and alerts

By default, all logs are sent to Priority Insights. Use the following criteria to decide which pipeline a log belongs to:

Criteria	Questions to ask	Recommended pipeline
Operational urgency	Do you need alerts within seconds or minutes?	Analyze and Alert
Troubleshooting need	Is this data needed for debugging but not real-time?	Analyze and Alert
Compliance retention	Is this log needed only for audit or compliance lookup?	Storage and Search

Volume and verbosity	Is this a high-volume debug or trace log?	Storage and Search
Business importance	Does it represent a business-critical transaction or error?	Priority Insights
Query frequency	How often do you query it? Hourly? Rarely?	Frequent: Priority Insights or Analyze and Alert; Rare: Storage and Search

Data pipeline selection criteria

Sending logs to IBM Cloud Logs

Send IBM Cloud platform application, infrastructure, and operational logs to your IBM Cloud Logs instance. Follow each section to ensure that you capture all Financial Services required logging.

Activity Tracker event routing (required)

IBM Cloud Activity Tracker captures auditable events in each region of IBM Cloud and is required for Financial Services compliance. The Activity Tracker Event Routing service allows you to send the collected events to IBM Cloud Logs for searching, alerting, and analysis.

After you provision IBM Cloud Logs, create an Activity Tracker event route:

1. Create an Activity Tracker event route that routes events to your IBM Cloud Logs instance.
2. Install the Activity Tracker extension for IBM Cloud Logs. The extension includes predefined dashboards, alerts, and rules.

For more information about searching and managing auditable events that are sent from Activity Tracker, see Activity Tracking events for IBM Cloud Logs.

Application logs (required)

Application logs are sent to IBM Cloud Logs by using the IBM Cloud Logs agent. The agent is installed on a compute resource where it can be configured to capture and send logs from running applications.

Start capturing application logs by first configuring your application name and subsystem name in IBM Cloud Logs. For more information, see Configuring the IBM Cloud Logs agent to set application and subsystem names.

After you configure the application and subsystem, follow the instructions for your compute platform:

Red Hat OpenShift on IBM Cloud

Install the IBM Cloud Logs agent on Red Hat OpenShift on IBM Cloud by using Helm.

Linux

Install the IBM Cloud Logs agent on Linux servers.

Windows

Install the IBM Cloud Logs agent on Windows servers.

IBM Cloud service logs (required)

IBM Cloud services automatically send logs to IBM Cloud Logs. However, some services offer additional logging features that are required for Financial Services compliance. If you use the services that are listed in the following table, follow the instructions for required configuration.

IBM Cloud service	Required configuration
IBM Cloud Databases for PostgreSQL	Enable PostgreSQL pgaudit.
IBM Cloud Databases for MongoDB	Provision with the Enterprise plan for audit logs to be sent to IBM Cloud Logs automatically.
IBM Cloud Databases for Elasticsearch	Provision with the Platinum plan for Kibana audit logging.

Red Hat OpenShift on IBM Cloud	Enable audit logging.

Required service logging configuration

IBM Cloud Logs extensions (optional)

Several extensions are available for IBM Cloud Logs that define alerts, rules, and dashboards for IBM Cloud services.

While not required, investigate whether any of the extensions would be useful for your use case. For more information, see the complete list of IBM Cloud Logs extensions.

Using the TCO Optimizer

With the IBM Cloud Logs TCO Optimizer, you can define how logs are distributed across the three data pipelines to keep log storage costs low. You create policies in the TCO Optimizer, and each policy determines which of the three data pipelines log data flows to based on the log's application, subsystem, and severity.

⚠ **Important:** You must have a Object Storage bucket provisioned and configured to archive your logs before you create a policy. For more information, see Configuring a bucket for IBM Cloud Logs.

Recommended TCO policies

To achieve the lowest possible cost for log storage by using the TCO Optimizer, create the following policies:

Policy name	Application	Subsystem	Severity	Priority
Storage and Search by default	All	All	All	Low
Analyze and Alert logs	One or more applications	All	All	Medium
Highest priority logs	One or more applications	One or more subsystems	Critical, Error	High

Recommended TCO Optimizer policies

With these three policies:

- All logs are sent to the Storage and Search data pipeline by default
- Logs that match the application that is defined in the Analyze and Alert logs policy are sent to the medium-priority Analyze and Alert data pipeline
- Logs that match the application or subsystem that is defined in the Highest priority logs policy with Critical or Error severity are sent to the Priority Insights data pipeline (most costly)

Accessing the TCO Optimizer

To access the TCO Optimizer:

1. Log in to your IBM Cloud account.
2. Click the navigation menu and select **Observability** > **Logging**.
3. Click the logging instance that you want to work with.
4. Click **Open Dashboard**.
5. In the dashboard navigation, click **Data Pipelines** > **TCO Optimizer**.

TCO Optimizer menu

Creating a TCO policy

The following example shows how to create a TCO Optimizer policy that sends all logs from the "OpenShift Marketplace" application to the Analyze and Alert data pipeline.

1. In the TCO Optimizer, click **Add new policy**.

Adding new TCO Optimizer policy

2. Enter a policy name, for example, "OpenShift Marketplace Policy".
3. Select the application, subsystem, and severity of the logs that the policy applies to. In this example, "OpenShift Marketplace" is selected as the application with all subsystems and all severities.

Selecting new TCO Optimizer policy attributes

4. Set the priority for the policy. The priority determines the pipeline for logs that are matched by the policy. In this example, select **Medium** to send OpenShift Marketplace logs to Analyze and Alert.

Selecting new policy priority

⚠ **Important:** Choosing **High** priority sends logs to Priority Insights, the most expensive data pipeline.

5. Click **Apply**.

Next steps

After you create the TCO Optimizer policies and configure the IBM Cloud Logs agent to supply application and subsystem names, your log data is sent to the appropriate IBM Cloud Logs data pipelines.

You can continue to configure IBM Cloud Logs to your specifications:

- Using Live Tail to tail logs
- Explore the IBM Cloud Logs dashboard and create your own custom dashboards
- Automatically set severity level metadata of log messages with the log agent
- Create and manage alerts
- Configure extensions in IBM Cloud Logs for monitoring IBM Cloud® Databases for MySQL or IBM Cloud Databases for PostgreSQL

Related information

- IBM Cloud Logs documentation
- IBM Cloud Logs FAQ
- IBM Cloud Logs API
- Alerting in IBM Cloud Logs
- Routing logs
- IBM Cloud Databases for PostgreSQL - Enable logging
- IBM Cloud for Financial Services

Compliance monitoring

You are required to continuously monitor for possible security flaws and changes in baseline configurations for which you should take corrective action. With Compliance Manager you can embed security checks into your every day workflows to help monitor for security and compliance. By monitoring for risks, you can identify security vulnerabilities and quickly work to mitigate the impact and fix the issue. By using Compliance Manager along with external integrations (such as, OpenShift Compliance Operator (OSCO), Tanium, NeuVector, and so on), you can build a robust approach for monitoring for security and compliance issues.

Using Compliance Manager

Compliance Manager provides a number of pre-defined profiles. Each profile is a collection of *controlsA technical, administrative, or physical safeguard designed to meet a set of defined security and privacy requirements. Controls exist to prevent, detect, or lessen the ability of a threat to exploit a vulnerability.*, and each control has one or more goals. Goals are pre-defined automated tests used to evaluate your posture against a control.

⚠ **Important:** Running a scan against a specific profile does not ensure regulatory compliance. The scan is intended to provide a point in time statement of your current posture for a specific group of resources.

The IBM Cloud for Financial Services profile provides a tailored set of goals that are mapped to the IBM Cloud Framework for Financial Services control requirements. This profile should always be used when leveraging the VPC reference architecture.

In addition, if you are using Red Hat OpenShift on IBM Cloud (whether in the VPC reference architecture or the Satellite reference architecture), then you

should leverage the Red Hat OpenShift on IBM Cloud Compliance Operator (OSCO) (OSCO) via the OSCO integration with SCC.

To start evaluating your resources, see the Getting started with Compliance Manager

Compliance best practices for Financial Services Validated services

The following table provides references to additional information for managing security and compliance for each service in the reference architecture.

Category	VPC reference architecture	Satellite reference architecture	Optional for both
Core	• VPC infrastructure services [1]	• Satellite	
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry	• Red Hat OpenShift on IBM Cloud [2] • Container Registry	
Networking	• VPC infrastructure services • Direct Link • Transit Gateway		
Storage	• Block Storage for VPC • Object Storage	• Object Storage	
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services	• App ID
Logging and monitoring	• Activity Tracker Event Routing • Compliance Manager • Flow Logs for VPC	• Activity Tracker Event Routing • Compliance Manager	
Integration			• Event Streams

For more information, see:

- Understanding compliance in IBM Cloud - provides details on compliance within platform services.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Change Management (CM)	CM-2 (2) Baseline Configuration | Automation Support for Accuracy and Currency CM-6 Configuration Settings CM-6 (1) Configuration Settings | Automated Management, Application, and Verification
Maintenance (MA)	MA-2 Controlled Maintenance

System and Information Integrity (SI)	SI-2 (2) Flaw Remediation | Automated Flaw Remediation Status SI-6 Security and Privacy Function Verification
Security Assessment and Authorization (CA)	CA-7 Continuous Monitoring

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Change Management (CM)	CM-2 (2) Baseline Configuration | Automation Support for Accuracy and Currency CM-6 Configuration Settings CM-6 (1) Configuration Settings | Automated Management, Application, and Verification
Maintenance (MA)	MA-2 Controlled Maintenance
System and Information Integrity (SI)	SI-2 (2) Flaw Remediation | Central Management SI-6 Security Functionality Verification
Security Assessment and Authorization (CA)	CA-7 Continuous Monitoring

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Operational monitoring

---

1. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
2. Satellite-enabled service which runs in your Satellite location. ↩

Operational monitoring

Operational monitoring for gauging system health is a important complement to monitoring for security and compliance. Proper operational monitoring can help you determine whether you need to fail over to an alternative storage or processing site. In addition, operational monitoring can help you determine whether operations have returned to normal after a system disruption. Operational metrics include measurements for CPU usage, memory usage, or API response times.

Since IBM Cloud Monitoring is not Financial Services Validated, the next two sections describe "bring your own" software solutions for Red Hat OpenShift on IBM Cloud and Virtual Servers for VPC.

⚠ **Important:** Generally speaking, you should strive to use only services which are Financial Services Validated in your solutions. However, depending on your circumstance there may be exceptions. See the best practice Use only services that are IBM Cloud for Financial Services Validated for more details and potential exceptions.

Red Hat OpenShift on IBM Cloud

You need to install your own software solution for monitoring in Red Hat OpenShift on IBM Cloud within your own VPC. There are various ways an operational monitoring solution can be implemented. See Setting up an operational monitoring solution for one example that uses Red Hat OpenShift on IBM Cloud Prometheus and Grafana.

Virtual Servers for VPC

You need to install your own software solution for monitoring in virtual server instances. See Setting up an operational monitoring solution for one example of sending virtual server instance metrics with Prometheus Node Exporter to Red Hat OpenShift on IBM Cloud Prometheus and Grafana.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control

Contingency Planning (CP)	CP-2 (3) Contingency Plan | Resume Mission and Business Functions CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-10 System Recovery and Reconstitution

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Contingency Planning (CP)	CP-2 (3) Contingency Plan | Resume Essential Missions / Business Functions CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-10 Information System Recovery and Reconstitution

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Data encryption at rest

Data encryption

Data encryption at rest

Protecting data against unauthorized disclosure, modification or destruction throughout the data lifecycle is of paramount importance in the IBM Cloud for Financial Services. Cryptographic controls must be in place in all regions and availability zones to protect the confidentiality and integrity of data. Effective key management enables you to keep control of your own master and root encryption keys, that protect the keys used for data at rest encryption.

Data encryption at rest in IBM Cloud

Within IBM Cloud, you must encrypt data at rest using encryption keys that you manage. For systems having a high security categorization you must use Hyper Protect Crypto Services to manage encryption keys. Hyper Protect Crypto Services is a single tenant key management service and dedicated hardware security module (HSM) partition. Hyper Protect Crypto Services is the only service in the cloud industry that's built on FIPS 140-2 Level 4-certified hardware.

With Hyper Protect Crypto Services you can take ownership of the cloud HSM to fully manage your encryption keys and to perform cryptographic operations by using the Keep Your Own Key (KYOK) concept. This means you have full control and authority over encryption keys. No one except you (not even IBM) has access to your encryption keys.

Some organizations require the use of a dedicated HSM where the device can be initialized by their cryptography specialists. We offer Hyper Protect Crypto Services with Bring Your Own HSM to enable those organizations to manage their HSM over the network at a remote location. This solution does uses the Hyper Protect Crypto Services key manager that has been Financial Services Validated but as IBM has no control over the Bring Your Own HSM, this solution as a whole isn't Financial Services Validated.

Where you have systems not requiring a high security categorization, you have the option of using IBM® Key Protect for IBM Cloud® for your key management. Key Protect is a multi-tenant cloud service using the Bring Your Own Key (BYOK) concept. Many capabilities are common between Hyper Protect Crypto Services and Key Protect but the underlying security implementation is different as shown in Table 1.

	Hyper Protect Crypto Services	**Hyper Protect Crypto Services with Bring Your Own HSM** [1]	**Key Protect**
Cloud REST API	X	X	X
Multi-Tenant Key Manager			X
Provider-managed HSM			X
Single-Tenant Key Manager	X	X	
Key Manager in Secure Enclave	X	X	
FIPS 140-2 Level 4 HSM	X		
Client-managed dedicated HSM partition	X		
HSM Smartcard key generation ceremony	X		
EP11 (PKCS#11) API	X		
MZR Availability	Selected	Selected	All
Support for Satellite location			X
UKO Key Generation [2]	X		

Key Features for IBM Cloud Key Management

Note 1: Capability depends on functionality of Bring Your Own HSM.

Hyper Protect Crypto Services and Key Protect are both Financial Services Validated with the same level of security assurance. The strength of the security in Hyper Protect Crypto Services, as a single-tenant service with smartcards for the master key ceremony, is higher than Key Protect as a multi-tenant service.

Whether you use Hyper Protect Crypto Services or Key Protect, you must ensure that:

- Keys are created for a single purpose, and dedicated encryption keys are unique per consumer.
- Keys have a lifecycle that is defined and are rotated periodically based on IBM Cloud Framework for Financial Services controls.
- Recovery functions, if used, are accessible only by authorized personnel.

Setting up Hyper Protect Crypto Services

For the provisioning and configuration of an instance of Hyper Protect Crypto Services see Getting started with Hyper Protect Crypto Services for specific instructions.

> **Tip:** For the highest level of security, you should use smart cards when initializing Hyper Protect Crypto Services. See Initializing service instances by using smart cards and the Management Utilities for more details.

For more background information and considerations, see:

- Components and concepts
- Bringing your encryption keys to the cloud
- Master key rotation
- Root key rotation

Continue with completing integration of Hyper Protect Crypto Services with the IBM Cloud services that are part of your deployment. Block Storage for VPC and Object Storage are common services that need data encrypted at rest, but there is an extended set of services that integrate. The Hyper Protect Crypto Services documentation includes a catalog of IBM Cloud services to integrate.

The master key within the Hyper Protect Crypto Services HSM and the root keys contained within Hyper Protect Crypto Services key manager need rotating with a new key at the frequency determined by your policy. If you don't have a policy, start with a 12 month rotation period. Consult the master key rotation process and root key rotation process for specific instructions.

For more background information and considerations, see:

- Components and concepts
- Bringing your encryption keys to the cloud

Setting up Hyper Protect Crypto Services with Bring Your Own HSM

As an alternative to using the in-built FIPS 140-2 Level 4 Cloud HSM for Hyper Protect Crypto Services, there is an option to Bring Your Own Hardware Security Module (BYOHSM). Details on the integration and the configuration is contained in the Introducing Bring Your Own HSM documentation.

The Bring Your Own HSM option provides you with the the benefit of being able to physically initialize your own HSM device. We understands for some organizations this is a mandatory requirement and this solution will meet that requirement.

However to operational characteristics of Hyper Protect Crypto Services changes. Here are a few things to keep in mind when selecting to use the Bring Your Own HSM option:

1. The Bring Your Own HSM may be FIPS 140-2 Level 3, which is a lower security level than the HSM used as default by Hyper Protect Crypto Services.
2. With the standard Hyper Protect Crypto Services solution there is negligible latency between the Hyper Protect Crypto Services key manager and Cloud HSM as they're within the same system. Ensure you assess any impact from increased latency between the HPCS key manager and your Bring Your Own HSM.
3. With the additional network hops between the Hyper Protect Crypto Services instance and the Bring Your Own HSM there may be a reduction in workload resilience. Review and understand the impact this change may have to your workload resiliency.
4. Ensure you identify an effective cold start sequence for HPCS connecting to your Bring Your Own HSM and the services that depend on Hyper Protect Crypto Services key management. For example, if the the connection of HPCS to the Bring Your Own HSM is through a firewall appliance using block storage with a key from Hyper Protect Crypto Services, you may get into a situation where the Bring Your Own HSM can't communicate from the firewall not being able to decrypt the block storage. This is otherwise known as the "keys locked in the car" scenario.

For more background information and considerations, see:

- Introducing Bring Your Own HSM
- Setting up Bring Your Own HSM
- Managing your keys with Bring Your Own HSM in IBM Cloud Hyper Protect Crypto Services

Setting up Key Protect

For the provisioning and configuration an instance of Key Protect, consult the Getting started tutorial for specific instructions.

Continue with completing integration of Key Protect with the IBM Cloud services that are part of your deployment. Block Storage for VPC and Object Storage are common services that need data encrypted at rest, but there is an extended set of services that integrate. The Key Protect documentation includes a catalog of IBM Cloud services to integrate.

The root keys contained within Key Protect need rotating with a new key at the frequency determined by your policy. If you don't have a policy, start with a 12 month rotation period. Consult the root key rotation process for specific instructions.

For more background information and considerations, see:

- About Key Protect
- Bringing your encryption keys to the cloud

Data encryption at rest in Satellite location

In an IBM Cloud Satellite location, the cloud services must use Key Protect on Satellite to connect to a user owned and managed hardware security module (HSM) for the encryption of data at rest. Consult About Key Protect on Satellite for specific information on the dedicated service.

⚠ **Important:** You are responsible for providing the underlying virtual and physical storage that will be accessed by Satellite storage in your on-premises Satellite location. You should ensure your underlying physical storage is encrypted using keys managed by an on-premises FIPS 140-2 level 2 or higher hardware security module (HSM).

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
System and Communications Protection (SC)	SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection SC-28 Protection of Information at Rest SC-28 (1) Protection of Information at Rest | Cryptographic Protection

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
System and Communications Protection (SC)	SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection SC-28 Protection of Information At Rest SC-28 (1) Protection of Information at Rest | Cryptographic Protection

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

In addition, you must follow the guidance found in Appendix A: Cryptographic Requirements within the "Control Implementation Overview Template for Application Providers Using IBM Cloud® Virtual Private Cloud".

Next steps

- Data encryption in transit

---

1. HSM capabilities depend on the technology provided for the Bring Your Own HSM component. ↩
2. UKO only supports Hyper Protect Crypto Services using an IBM Cloud HSM. ↩

Data encryption in transit

Data in transit must always be encrypted. Cryptographic controls must be in place in all regions and availability zones to protect the confidentiality and integrity of data. Data must be protected by at least two levels of encryption such as transport and field level or envelope level.

As an application provider, you must:

- Ensure that all traffic, tunneling, and other protection of data in transit is encrypted by using TLS 1.2 at a minimum. You must disable the SSL/TLS 1.0/TLS 1.1 protocols on all software that is deployed in your VPCs. See the following TLS requirements section for more details.
- Use Hyper Protect Crypto Services for TLS offload for all data that is requested from outside of IBM Cloud (inbound traffic to IBM Cloud) and protected by a certificate that is signed by a public certificate authority. In this scenario, you should configure any web servers to use TLS offload to set up the session so that the private key never leaves Hyper Protect Crypto Services. See Use IBM Cloud Hyper Protect Crypto Services to offload NGINX TLS for an example.
- Enforce encryption that uses directives like HTTP Strict Transport Security (HSTS).
- Disable caching for responses that contain sensitive data.
- Never send confidential information with URL parameters.
- Use HTTPS for all connections from the web. HTTP connections must be refused or redirected to HTTPS.

⚠ **Important:** The requirement to encrypt data in transit also applies to traffic within your VPCs. This might be traffic from one virtual server instance to another, or from one Red Hat OpenShift on IBM Cloud container to another. If using containers, you might consider Red Hat OpenShift on IBM Cloud Service Mesh which can enable mTLS between containers without code changes.

TLS requirements

To be effective, TLS must be configured correctly otherwise the data in transit is vulnerable to modification or eavesdropping. The following steps are required to deploy a secure TLS connection:

- Securely generate and store a private key.
- Obtain a certificate associated with the private key, signed by the appropriate trusted certificate authority (CA).
- Configure the TLS server to use approved TLS versions (1.2 or 1.3) and approved cipher suites.
- Configure the TLS client to check certificate validity and use approved TLS versions only.
- Deploy certificate and private key to TLS server.
- Enable the TLS server.
- Replace the certificate before it expires (typically certificates expire after 1 year).

TLS protocol requirements

1. Only TLS versions 1.2 or 1.3 are permitted. All older versions of TLS and all version of SSL must not be used.
2. Prohibited versions of SSL/TLS must be disabled on the server to prevent downgrading to a vulnerable version of the protocol.
3. Only cipher suites that are listed under Acceptable cipher suites can be enabled. All others must be disabled.

Key size requirements

Key size requirements are as follows:

- RSA – minimum key size is 2048, and 4096 is recommended.
- ECDSA – minimum key size is 224.
- ECDHE – minimum key size is 224.

Acceptable cipher suites

For TLS 1.2, the acceptable cipher suites are as follows:

- TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256
- TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256

For TLS 1.2, the acceptable cipher suites are as follows:

- TLS_AES_128_GCM_SHA256

- TLS_AES_256_GCM_SHA384
- TLS_CHACHA20_POLY1305_SHA256

Certificate requirements

1. Certificate-private keys must be unique to that certificate. Private keys must not be reused for any other purpose
2. Certificate-private keys must not be reused when a certificate is renewed.
3. Certificates must specify the FQDN for the certificate in the **Subject Alternative Name** field; single-domain certificates should have that domain repeated in the SAN field.
4. TLS certificates must be used only as specified by the Enhanced Key Usage attribute within the certificate. The following key uses are acceptable for TLS certificates: TLS Web Server Authentication, TLS Web Client Authentication
5. External facing certificates that are issued by an external CA must have a maximum lifetime of 1 year.
6. Private keys of external facing certificates must be stored in HSM
7. Certificates should uniquely identify an endpoint and any subject alternative names, rather than using wildcards. Wildcard certificates should be used only when TLS endpoints are programmatically generated and must be limited to the smallest usable subdomain.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
System and Communications Protection (SC)	SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
System and Communications Protection (SC)	SC-8 Transmission Confidentiality and Integrity SC-8 (1) Transmission Confidentiality and Integrity | Cryptographic Protection SC-12 Cryptographic Key Establishment and Management SC-12 (2) Cryptographic Key Establishment and Management | Symmetric Keys SC-12 (3) Cryptographic Key Establishment and Management | Asymmetric Keys SC-13 Cryptographic Protection

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

In addition, you must follow the guidance found in Appendix A: Cryptographic Requirements within the Control Implementation Overview Template for Application Providers Using IBM Cloud® Virtual Private Cloud.

Next steps

- Business continuity and disaster recovery

Business continuity and disaster recovery

Business continuity and disaster recovery overview

Business continuity and disaster recovery (BCDR) is important for all cloud-based applications, and it is all the more critical for the kind of regulated workloads that the IBM Cloud for Financial Services is intended to host. In the event of a disaster, consumers depend on service being restored quickly without loss of data. Using the VPC reference architecture and following the controls of the IBM Cloud Framework for Financial Services help you achieve that goal. When building a BCDR solution in IBM Cloud, you need to consider your workloads that run in virtual servers or Red Hat OpenShift on IBM Cloud and the BCDR characteristics of all of the other IBM Cloud services your solution depends on.

Requirements

The following sections describe some of the requirements that you must follow for BCDR of your workloads.

Alternative storage site

You must establish and configure an alternative storage site in at least one geographically separate IBM Cloud multizone region. So, if the storage in one region becomes unavailable, you can use storage from another region.

Information system backup

As the provider, you must:

- Conduct backups of user-level information contained in the information system
- Conduct backups of system-level information contained in the information system
- Conduct backups of information system documentation, including security-related documentation
- Protect the confidentiality, integrity, and availability of backup information at storage locations

The different levels of information are defined as follows:

- User-level information - Consumer/consumer-managed data
- System-level information - Data you manage that is necessary to meet your service's Recovery Point Objective (RPO). For example, system-state information, operating system and application software, licenses, and so on.
- System documentation - External documentation, internal architecture documentation, run books, and so on.

Other requirements:

- Backups must be at least incremental daily and full weekly.
- Backups must be monitored. Failed backups must be investigated and corrective action that is documented as necessary to maintain the service's RPO.
- Backups and backup logs should be accessible only to authorized individuals who need access to do their jobs.
- Transaction-based systems must implement transaction recovery. Transaction-based systems include database management systems and transaction processing systems. Mechanisms supporting transaction recovery include transaction rollback and transaction journaling.

Information system recovery and reconstitution

You should enable recovery and reconstitution of the system to a known state after a disruption, compromise, or failure. Contingency plans must include procedures for validating successful recovery and reconstitution. Recovery and reconstitution activities include resuming operational capabilities at the original location.

Backup and disaster recovery for the reference architectures

See the following resources depending on which reference architecture you are using:

- Business continuity and disaster recovery for VPC reference architecture
- Business continuity and disaster recovery for Satellite reference architecture

Backup and disaster recovery for IBM Cloud services

In addition to backing up your workloads and having the capability to restore service in the face of a disaster, your BCDR strategy needs to consider all of the IBM Cloud services in your deployment. The following table provides references to additional information regarding BCDR for each service in the reference architecture.

The following table provides references for more information about BCDR for each service in the reference architecture.

Category	VPC reference architecture	Satellite reference architecture	Optional for both
Core	• VPC infrastructure services [1]	• Satellite	
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry	• Red Hat OpenShift on IBM Cloud [2] • Container Registry	
Networking	• VPC infrastructure services • Direct Link • Transit Gateway		
Storage	• Block Storage for VPC • Object Storage	• Object Storage	
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services	• App ID
Logging and monitoring	• Activity Tracker Event Routing • Compliance Manager • Flow Logs for VPC	• Activity Tracker Event Routing • Compliance Manager	
Integration			• Event Streams

Backup and disaster recovery information for IBM Cloud services in the reference architectures

In addition to the Financial Services Validated services in the reference architecture, see the following references for other important information:

- How IBM Cloud ensures high availability and disaster recovery
- High availability in the VPC reference architecture
- Responsibilities for operating services in your deployment of the VPC reference architecture

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
Contingency Planning (CP)	CP-2 Contingency Plan CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-9 System Backup CP-10 System Recovery and Reconstitution CP-10 (2) Information System Recovery and Reconstitution | Transaction Recovery

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control

Contingency Planning (CP)	CP-2 Contingency Plan CP-6 Alternate Storage Site CP-7 Alternate Processing Site CP-9 Information System Backup CP-10 Information System Recovery and Reconstitution CP-10 (2) System Recovery and Reconstitution | Transaction Recovery

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Development processes and software integrity

---

1. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
2. Satellite-enabled service which runs in your Satellite location. ↩

Business continuity and disaster recovery for VPC reference architecture

As described in Business continuity and disaster recovery (BCDR) overview, BCDR is important for all cloud-based applications. In this article, you will learn about specifics of BCDR in the reference architecture.

Workloads in VPC reference architecture

Workloads on virtual server instances

If you are running virtual server instances, then you can consider Snapshots for VPC. Snapshots for VPC is a regional offering that lets you create a point-in-time copy of your block storage boot or data volumes. The initial snapshot you take is a full backup of the volume. Subsequent snapshots of the same volume are incremental; only the changes since the last snapshot are captured. You can select a snapshot during instance provisioning, and restore a new, fully-provisioned boot volume to start the instance. You can also create and attach a data volume from a snapshot within a running virtual server instance.

For more information, see General procedure for creating and using snapshots.

Workloads on Red Hat OpenShift on IBM Cloud

If you are running Red Hat OpenShift on IBM Cloud, then you can consider Portworx. Portworx is a highly available software-defined storage solution that you can use to manage local persistent storage for your containerized databases and other stateful apps, or to share data between pods across multiple zones.

See Storing data on software-defined storage (SDS) with Portworx for more information.

Backup and disaster recovery for IBM Cloud services

In addition to backing up your workloads and having the capability to restore service in the face of a disaster, your BCDR strategy needs to consider all of the IBM Cloud services in your deployment. See Backup and disaster recovery for IBM Cloud services.

Related controls in IBM Cloud Framework for Financial Services

See related controls for backup and recovery.

Next steps

- High availability overview

Business continuity and disaster recovery for Satellite reference architecture

As described in Business continuity and disaster recovery (BCDR) overview, BCDR is important for all cloud-based applications. You will learn about specifics of BCDR in the Satellite reference architecture.

BCDR in your Satellite location

See the Disaster recovery section of Satellite for more information.

Backup and disaster recovery for IBM Cloud services

In addition to backing up your workloads and being able to restore service in the face of a disaster, your BCDR strategy needs to consider all of the IBM Cloud services in your deployment. See Backup and disaster recovery for IBM Cloud services .

Related controls in IBM Cloud Framework for Financial Services

See related controls for backup and recovery .

Next steps

- High availability overview

High availability

High availability overview

Ensuring high availability (HA) is important for all cloud-based applications, and it is all the more critical for regulated workloads that the IBM Cloud for Financial Services is intended to host. In short, consumers depend on your application and cannot afford downtime. Using the VPC reference architecture and following the controls of the IBM Cloud Framework for Financial Services help you achieve that goal. When building an HA solution in IBM Cloud, you need to consider your workloads that run in virtual servers or Red Hat OpenShift on IBM Cloud and the HA characteristics of all of the other IBM Cloud services your solution depends on.

HA for the reference architectures

See the following resources depending on which reference architecture you are using:

- High availability for VPC reference architecture
- High availability for Satellite reference architecture

HA in IBM Cloud services

Your HA strategy needs to consider all of the IBM Cloud services in your deployment. The following table provides references to additional information on HA for each service in the reference architecture.

Category	VPC reference architecture	Satellite reference architecture	Optional for both
Core	• VPC infrastructure services [1]	• Satellite	
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry	• Red Hat OpenShift on IBM Cloud [2] • Container Registry	
Networking	• VPC infrastructure services • Direct Link • Transit Gateway		
Storage	• Block Storage for VPC • Object Storage	• Object Storage	
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services	• App ID
Logging and monitoring	• Activity Tracker Event Routing • Compliance Manager • Flow Logs for VPC	• Activity Tracker Event Routing • Compliance Manager	
Integration			• Event Streams

High availability information for IBM Cloud services in the reference architectures

The references cover the steps that the service has taken to achieve high availability as well as the steps you must take. For example, consider Direct Link. The Direct Link service itself makes use of multiple zones within a region to help achieve HA. However, you Direct Link is not a redundant service at the cross-connect router (XCR), so you have the responsibility for creating redundancy through your border gateway protocol (BGP) schemas. See High availability and disaster recovery for Direct Link and Models for diversity and redundancy in Direct Link for more details.

In addition to the Financial Services Validated services in the reference architecture, see the following references for other important information:

- How IBM Cloud ensures high availability and disaster recovery
- Responsibilities for operating services in your deployment

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, ydo your own due diligence to ensure you meet the requirements.

Family	Control
Contingency Planning (CP)	CP-2 Contingency Plan CP-7 Alternate Processing Site
System and Communications Protection (SC)	SC-6 Resource Availability

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
Contingency Planning (CP)	CP-2 Contingency Plan CP-7 Alternate Processing Site
System and Communications Protection (SC)	SC-6 Resource Availability

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Development processes and software integrity

---

1. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
2. Satellite-enabled service which runs in your Satellite location. ↩

High availability for VPC reference architecture

The High availability (HA) overview desribes how HA is important for all cloud-based applications. In this article, you will learn more specifics about HA in the VPC reference architecture.

Your workloads on VPC infrastructure

Workloads in multizone regions

Distribute workloads across zones

Recall that all VPCs exist in a single region referring to the geographic area in which a VPC is deployed. It is recommended that you distribute your workloads across multiple zones within a region as the first step in creating an HA solution. That is, if one zone goes down, the other zones can continue to support the workload. You should use Application Load Balancer for VPC (ALB) to distribute your traffic across zones in a region.

Scale based on resource demands

Even with multiple zones, you should take additional steps to ensure that your environment can be scaled out based on the resources used by your workloads to enhance resiliency.

Virtual servers

With Auto Scale for VPC, you can improve performance and costs by dynamically creating virtual server instances to meet the demands of your environment. Auto Scale for VPC is highly recommended if you are using virtual servers. You set scaling policies that define your desired average utilization for metrics like CPU, memory, and network usage. The policies that you define determine when virtual server instances are added or removed from your instance group. Auto Scale for VPC is highly recommended if you are using virtual servers.

Containers

For Red Hat OpenShift on IBM Cloud, you can use the `cluster-autoscaler` add-on to scale the worker pools in your cluster automatically to increase or decrease the number of worker nodes in the worker pool based on the sizing needs of your scheduled workloads. The `cluster-autoscaler` add-on is based on the Kubernetes Cluster-Autoscaler project. See Autoscaling clusters for more information.

In addition, it is possible to autoscale your pods. See Scaling apps for more information.

Workloads in multiple regions

Despite the steps you take to leverage multiple zones in a region and increase resiliency with auto scaling, it is still possible for an entire region to be taken out of service. For example, a natural disaster like a hurricane, tornado, or earthquake could knock out multiple zones in a region.

So, it is also recommended to establish and configure an alternate processing site in at least one geographically separate multizone region to ensure the continuation of secure system operation. Then, in addition to spreading your workloads across zones in a region, you can distribute your workloads across regions. However, your consumers must be able to opt out of having their data stored in the alternate region based on their data residency requirements.

The diagram below shows an expansion of the VPC reference architecture to use multiple regions. The key to making this work is to use the global load balancing functionality in DNS Services. The global load balancer offers HA and geographical distribution of your traffic, based on the health of your origin servers and the geographical region where the user request originates. So, if one region becomes unhealthy, then traffic would get routed to the next closest healthy region.

VPC architecture in multiple regions

For completeness, the diagram below shows the VPC architecture spread across multiple regions when Red Hat OpenShift on IBM Cloud is used.

Red Hat OpenShift on IBM Cloud on VPC multiple regions

See the following tutorials for more information:

- Strategies for resilient applications
- Deploy isolated workloads across multiple locations and zones

HA for IBM Cloud services

Your high availability strategy needs to consider all of the IBM Cloud services in your deployment. See High availabiltty for IBM Cloud services for more information.

Related controls in IBM Cloud Framework for Financial Services

See related controls for high availability.

Next steps

- Development processes and software integrity

High availability for Satellite reference architecture

As described in High availability (HA) overview, HA is important for all cloud-based applications. Ready to learn more specifics about HA in the Satellite reference architecture?

High availability in your Satellite location

For HA in IBM Cloud Satellite, you need to consider HA in three main areas within your Satellite location. For more information, see the following topics:

1. Satellite control plane master
2. Satellite control plane worker nodes
3. IBM Cloud services that run in your Satellite location such as Red Hat OpenShift on IBM Cloud

In addition, you own all of the underlying infrastructure that the Satellite location components run on. Consider the following options to increase resiliency of that infrastructure:

- Add more compute hosts
- Add redundant power, network, and storage
- Isolate machines within one data center
- Spread hosts across physical locations

Your workloads in Red Hat OpenShift on IBM Cloud

After you have your underlying infrastructure in place, begin deploying your workloads in your on-premises Red Hat OpenShift on IBM Cloud workload clusters. You can use the `cluster-autoscaler` add-on to scale the worker pools in your cluster automatically to increase or decrease the number of worker nodes in the worker pool based on the sizing needs of your scheduled workloads. The `cluster-autoscaler` add-on is based on the Kubernetes Cluster-Autoscaler project. For more information, see Autoscaling clusters.

In addition, it is possible to autoscale your pods. For more information, see Scaling apps.

HA for IBM Cloud services

Your high availability strategy needs to consider all of the IBM Cloud services in your deployment, including the services that run in IBM Cloud outside of your Satellite location. For more information, see High availability for IBM Cloud services.

Related controls in IBM Cloud Framework for Financial Services

For more information, see related controls for high availability.

Next steps

- Development processes and software integrity

Development processes and software integrity

Development processes and ensuring system and software integrity are key parts of the IBM Cloud Framework for Financial Services. You must manage changes to the information system by using a defined system development life cycle (SDLC) that incorporates information security considerations.

You should ensure that security engineering principles are applied in the design, development, implementation, and modification of the system. During the ongoing development, implementation, operations, maintenance, and continuous monitoring of the overarching information system lifecycle, you must develop, implement, operate, and maintain the system in accordance with the security processes and technical requirements that are detailed in the IBM Cloud Framework for Financial Services and the accompanying operational plans and policies.

You must ensure that developers:

- Perform configuration management during system/component/service development, implementation, and operation.
- Document change approvals and security impact analyses for all system changes.
- Perform unit, integration, system, regression testing/evaluation.
- Perform static and dynamic code analysis.
- Use development testing processes that provide coverage for the entire component and explicitly reviews, evaluates, and tests all security functions.
- Control consumer data so that it is never placed into nonproduction environments. Consumer data must not be consumed or used for the purposes of testing services.
- Enable integrity verification of software components to detect unauthorized changes to system software, firmware, and information:
  - All applicable software components must be cryptographically signed by the manufacturer or developer to ensure you can perform integrity checks and verify that the component that is deployed in the system is the same component that the manufacturer or developer evaluated and certified.
  - Deploy only signed and approved builds within the environment.
- Ensure only the approved, tested changes/code gets promoted to production and no vulnerabilities are introduced.
- Document and evidence the execution of the system/service and security testing/scanning along with the results.
- Track security flaws and flaw resolution within the system and report findings to designated personnel.

Red Hat OpenShift on IBM Cloud

When using Red Hat OpenShift on IBM Cloud to host workloads, you must use Container Registry and the Vulnerability Advisor component it contains. Red Hat OpenShift on IBM Cloud provides a multi-tenant, highly available, scalable, and encrypted private image registry that is hosted and managed by IBM®. You can use Container Registry by setting up your own image namespace and pushing container images to your namespace. By using Container Registry, only users with access to your IBM Cloud account can access your images.

When you push images to Container Registry, you benefit from the built-in Vulnerability Advisor features that scan for potential security issues and vulnerabilities. Vulnerability Advisor checks for vulnerable packages in specific Docker base images, and known vulnerabilities in app configuration settings. When vulnerabilities are found, information about the vulnerability is provided. You can use this information to resolve security issues so that containers are not deployed from vulnerable images.

Any issues that are found by Vulnerability Advisor result in a verdict that indicates that it is not advisable to deploy this image. If you choose to deploy the image, any containers that are deployed from the image include known issues that might be used to attack or otherwise compromise the container. The verdict is adjusted based on any exemptions that you specified. This verdict can be used by Portieris to prevent the deployment of nonsecure images in Container Registry. Portieris is a Kubernetes admission controller for the enforcement of image security policies. You can create image security policies for each Kubernetes namespace, or at the cluster level, and enforce different rules for different images.

Fixing the security and configuration issues that are reported by Vulnerability Advisor can help you to secure your IBM Cloud infrastructure.

See the following for more information on Container Registry and how to set it up:

- About Container Registry
- Setting up Container Registry as a private registry on Red Hat OpenShift on IBM Cloud
- Managing image security with Vulnerability Advisor
- Signing images for trusted content
- Container Registry and Vulnerability Advisor workflow tutorial

Virtual server instances

⚠ **Important:** There is not a Financial Services Validated solution for scanning virtual server images. You need to install your own software solution.

Related controls in IBM Cloud Framework for Financial Services

The following IBM Cloud Framework for Financial Services controls are most related to this guidance. However, in addition to following the guidance here, do your own due diligence to ensure you meet the requirements.

Family	Control
System and Information Integrity (SI)	SI-7 Software, Firmware, and Information Integrity
System and Services Acquisition (SA)	SA-3 System Development Life Cycle SA-3 (2) System Development Life Cycle | Use of Live or Operational Data SA-8 Security and Privacy Engineering Principles SA-10 Developer Configuration Management SA-10 (1) Developer Configuration Management | Software and Firmware Integrity Verification SA-11 Developer Testing and Evaluation SA-15 Development Process, Standards, and Tools
Risk Assessment (RA)	RA-5 Vulnerability Monitoring and Scanning

Related controls in IBM Cloud Framework for Financial Services [FSv2.0]

Family	Control
System and Information Integrity (SI)	SI-7 Software & Information Integrity
System and Services Acquisition (SA)	SA-3 System Development Life Cycle SA-8 Security Engineering Principles SA-10 Developer Configuration Management SA-10 (1) Developer Configuration Management | Software and Firmware Integrity Verification SA-11 Developer Security Testing and Evaluation SA-15 Development Process, Standards, and Tools SA-15 (9) Development Process, Standards, and Tools | Use of Live Data
Risk Assessment (RA)	RA-5 Vulnerability Scanning

Related controls in IBM Cloud Framework for Financial Services [FSv1.1]

Next steps

- Responsibilities for operating services in your deployment

Responsibilities for operating services in your deployment

In IBM Cloud, the responsibilities for deploying, operating, and securing products are shared between IBM and you, the application provider. It is important for you to understand these responsibilities for every IBM Cloud service you've deployed as part of the reference architectures.

Depending on the product, the responsibility for the following types of tasks (which span activities for Day 0, Day 1, and Day 2) can be exclusive to IBM, you, or shared. The tasks for each type of product are grouped in the following categories, which cut across nearly all of the IBM Cloud Framework for Financial Services's best practices and requirements.

Incident and operations management

: Includes tasks such as monitoring, event management, high availability, problem determination, recovery, and full state backup and recovery.

Change management

: Includes tasks such as deployment, configuration, upgrades, patching, configuration changes, and deletion.

Identity and access management

: Includes tasks such as authentication, authorization, access control policies, and approving, granting, and revoking access.

Security and regulation compliance

: Includes tasks such as security controls implementation and compliance certification.

Disaster recovery

: Includes tasks such as providing dependencies on disaster recovery sites, provision disaster recovery environments, data and configuration backup, replicating data and configuration to the disaster recovery environment, and failover on disaster events.

For more information, see Shared responsibilities for using IBM Cloud products.

Responsibilities for Financial Services Validated services

The following table provides references that describe your responsibilities for each Financial Services Validated service in the reference architecture.

Category	VPC reference architecture	Satellite reference architecture	Optional for both
Core	• VPC infrastructure services [1]	• Satellite	
Containers	• Red Hat OpenShift on IBM Cloud • Container Registry [2]	• Red Hat OpenShift on IBM Cloud [3] • Container Registry	
Networking	• VPC infrastructure services • Direct Link • Transit Gateway		
Storage	• Block Storage for VPC • Object Storage	• Object Storage	
Security	• Hyper Protect Crypto Services	• Hyper Protect Crypto Services	• App ID [4]

Logging and monitoring	• Activity Tracker Event Routing • Compliance Manager • Flow Logs for VPC	• Activity Tracker Event Routing • Compliance Manager
Integration		• Event Streams

Your responsibilities for IBM Cloud services in the reference architectures

1. Includes VPC, Dedicated hosts for VPC, Auto Scale for VPC, Application Load Balancer for VPC, VPN for VPC, DNS Services, and VPE for VPC. ↩
2. There is not a specific shared responsibility article for Container Registry. ↩
3. Satellite-enabled service which runs in your Satellite location. ↩
4. There is not a specific shared responsibility article for App ID. ↩

IBM Cloud for Financial Services Validation Process for Partners

IBM designates ecosystem partners' offerings as IBM Cloud for Financial Services Validated when the partners' offerings conform to the IBM Cloud for Financial Services® reference architectures and guidance and have been determined by IBM to materially implement the IBM Cloud Framework for Financial Services control requirements.

The IBM ISV/MSP Acceleration Program works with partners through the process, accelerating time to commercial readiness and reducing the partners' challenge of providing their solution to IBM Cloud customers.

The Financial Services Validation Process for partners follows the following process:

1. Initial meeting
    - Partner and IBM discuss benefits of participating in IBM Cloud for Financial Services ecosystem
    - Verify current NDA in place before discussing confidential information
    - ISV identifies the application(s) and workload(s) they will bring
    - Qualifying questions and discussion about client demand/path to revenue
2. Onboarding kickoff meeting is held
    - Walkthrough of detailed onboarding process steps
    - Service technology stack (Cloud Native: Containers: ROKS, Cloud Native: Virtual Servers, VMWare) decision
    - Service delivery model: BYOL Services, BYOL Software, Hosted Managed Service, Joint Delivery, SaaS: Multi-tenant, SaaS: Single-tenant) decision
    - ISV begins technical and architecture assessments ensuring they can conform to architecture requirements
3. ISV onboards their service to the IBM ecosystem
    - ISV becomes a member of Partner Plus (IBM Ecosystem partnership program)
    - Implementation of IBM Cloud for Financial Services controls framework for the ISV solution
    - ISV signs an ESA (Embedded Solution Agreement) if IBM Cloud is embedded in their solution
    - Optionally work to list the ISV solution in the IBM Cloud Catalog (with ESA)
4. ISV completes Security Assessments and demonstrates ability to meet applicable IBM Cloud for Financial Services controls and provides evidence
    - Based on controls specific to Software and/or SaaS model
    - Security & Compliance Center configured
5. IBM Validation Team completes testing and review to designate ISV offering “Financial Services Validated”
    - ISV enabled to share CIOD report with customers as applicable

Notes:

- Some customers may be willing to conduct POC on IBM Cloud with the partner offering before FS Validation is complete
- Periodic reassessment required to maintain designation

To onboard to the IBM Cloud catalog, the following guidance is available:

- Getting set up to sell services
- Getting set up to sell software

IBM Cloud for Financial Services Validated Partner Offerings

IBM designates ecosystem partners' offerings as IBM Cloud for Financial Services Validated when the partners' offerings conform to the IBM Cloud for Financial Services® reference architectures and guidance and have been determined by IBM to materially implement the IBM Cloud Framework for Financial Services control requirements. Clients should review associated documentation made available by IBM and ecosystem partners to conduct their due diligence.

Through the shared responsibility model of the IBM Cloud Framework for Financial Services and the surrounding standardized processes, financial institutions and ecosystem partners get benefits such as:

- Less time that is spent by ecosystem partners demonstrating compliance and more time delivering innovative services.
- Reduction in the time and effort by financial institutions to ensure the compliance of third-party vendors, and more time spent delivering new, innovative services to their customers.
- Streamlined procurement, contracting, and onboarding within the ecosystem that leads to reduced time to market for all parties.

Ecosystem partners (even those who are not yet IBM Cloud for Financial Services Validated) are encouraged to onboard to the IBM Cloud catalog.

List of FS Validated ISV Offerings

> Note this list only includes partners who have completed the FS validation process with IBM and have been evaluated by IBM to have correctly implemented the Framework. Partners going through the process now are not listed. There are also many more ISVs and Partners not listed here that make up the broader ecosystem of IBM Cloud for Financial Services.

Contact IBM Sales or open a support case if you are interested in any of these offerings, or if you don't see an offering you are interested in on the list.

ISV	Offering
Adobe	Adobe Experience Manager
Appian	Appian Platform
Circeo	The Loan Factory
Dizzion	Dizzion Complete FS - Horizon
Effyis	P2S Platform
EIS	EIS Suite
EpositBox	EpositBox
F5	F5 BIG-IP
Gambit	Wealth as a Service
InvestSuite	Wealth as a Service
Kanexa/Marco Polo Network	Supplier Pay
Liferay	DXP
NexJ SYSTEMS	Customer Relationship Management for Wealth Management
Palo Alto networks	Prisma Cloud Compute Edition, VM-Series Firewall
Portworx	PX-Backup for Kubernetes, Portworx Enterprise
QuickSign	QuickSign

Quid	Qinetic
Ridecell	Rideval Fleet Optimizer
Ripjar	Labyrinth
SAP	SAP S/4 HANA, ASE (Sybase), SAP RISE
SAS	Viya
ServiceNow	ServiceNow Managed Service
Sopra Banking Software	Sopra Banking Software for Lending, CMS Card Management Solution
Tagetik	CCH Tagetik Intelligent Platform
Tanium	Tanium Cloud Appliance
Temenos	Data Hub, Financial Crime Mitigation (FCM), Payment Hub, Transact, WealthSuite
Topaz	Core Banking
Tovie.ai	Conversational AI Solutions
Unblu	Unblu
Upsun	Upsun Fixed
Veeam	Veeam Availability Suite
WorkFusion	Intelligent Automation of Customer Lifecycle
Yields.io	Chiron/MRM

FAQs for IBM Cloud for Financial Services

Answers to common questions about IBM Cloud for Financial Services and the IBM Cloud Framework for Financial Services.

What are the supported reference architectures for the IBM Cloud for Financial Services?

See references architectures.

What is a *consumer* and what is an *application provider*?

We refer to the entity that uses (consumes) the application as the *consumer* and the entity that deploys (provides) the application as the *application provider*. This documentation generally assumes that you are a an application provider.

The application provider and consumer might be the same when the organization providing the application also consumes it (for example, a financial institution deploys an application for internal use by others in their company). Or, they can be different when a third-party ISV provides an application for a financial institution to use.

Application providers who are within financial institutions are not beholden to the IBM Cloud Framework for Financial Services control requirements.

What does it mean for a service to be IBM Cloud for Financial Services Validated?

IBM Cloud for Financial Services Validated designates that an IBM Cloud service or ecosystem partner service has evidenced compliance to the controls of the IBM Cloud Framework for Financial Services and can be used to build solutions that might themselves be validated.

Should I only use IBM or third-party managed services that are IBM Cloud for Financial Services Validated?

Generally speaking, you should strive to use only services which are Financial Services Validated in your solutions. However, depending on your circumstance there may be exceptions. See the best practice Use only services that are IBM Cloud for Financial Services Validated for more details and potential exceptions.

What if I don't want to use managed services, but I want to install and manage my own software?

Self-installed (or "bring your own") software from third parties is permitted without special approval. Examples include databases, logging stacks, web application firewalls, and so on. But just like first-party software you might develop and deploy, self-installed third-party software must comply with all of the requirements of the IBM Cloud Framework for Financial Services. And, it is your responsibility to provide appropriate supporting evidence.

For more information, see Ensure all usage of software meets IBM Cloud Framework for Financial Services controls.

Site map

Find what you are looking for in the topics that are available for IBM Cloud for Financial Services®.

Getting started with IBM Cloud for Financial Services

Security categorization, functionality and assurance

Best practices and requirements

- 15. Regularly scan for open ports / protocols
- 16. Secure and manage secrets and certificates
- 17. Tag all IBM Cloud resources with security attributes
- 18. Monitor for security and compliance against a baseline configuration
- Next steps

Best practices and requirements for software

- 1. Enable operators to properly configure your application to meet the controls
- 2. Provide account and access management to enable a zero-trust environment
- 3. Enable ability to capture audit events
- 4. Follow secure development processes and ensure software integrity
- 5. Ensure consumer data can be encrypted at rest and in transit
- 6. Enable software to be run in a highly available manner

Release notes

Release notes

- 23 June 2026
- 17 June 2026
- 15 June 2026
- 10 June 2026
- 2 June 2026
- 24 March 2025
- 30 April 2023
- 24 April 2023
- 31 March 2023
- 03 March 2023
- 23 January 2023
- 30 September 2022
- 24 August 2022
- 28 July 2022
- 17 March 2022
- 03 February 2022
- 14 October 2021
- 07 April 2021

Reference architecture overview

Reference architecture overview

IBM Virtual Private Cloud (VPC) architecture

VPC reference architecture

- Next steps

VPC virtual servers and Power virtual servers reference architecture

- Architecture diagram
- Design concepts
- Requirements
- Components

VPC with Red Hat OpenShift on IBM Cloud reference architecture

- Architecture diagram
- Red Hat OpenShift on IBM Cloud concepts
- Next steps

IBM Cloud Satellite architecture

Satellite reference architecture

- Solution design
- Financial Services Validated services
  - Core
  - Containers
  - Storage
  - Security
  - Logging and monitoring
  - Integration
- Satellite components in IBM Cloud
- Satellite location components
  - Infrastructure
  - Satellite control plane worker nodes
  - Workload clusters
- Other supporting components on-premises
- Next steps

Satellite reference architecture use cases

- US regional bank
- Multinational bank
- Next steps

Satellite reference architecture shared responsibility model

- Shared responsibilities for IBM Cloud products
- Additional roles for shared responsibility matrix
- Overview of shared responsibilities
- Satellite location control plane and workload clusters
- Other on-premises components outside of the IBM Cloud Satellite location
- Next steps

Satellite reference architecture best practices checklist

- Relation to IBM Cloud Framework for Financial Services best practices
- Next steps

IBM Cloud for VMware Regulated Workloads architecture

IBM Cloud for VMware Regulated Workloads architecture

Bring your own solutions for VPC reference architecture samples

Web application firewall and full-tunnel VPN

Deploying and configuring F5 BIG-IP

- Objectives
- Before you begin
- Provision BIG-IP
- Complete licensing for BIG-IP
- Configure management UI cipher suite
- Replacing the configuration utility's self-signed device certificate
- Accessing the F5 BIG-IP
- Enabling GUI audit logging
- Setup BIG-IP routes
- Access control lists and security groups
  - F5 management interface
  - F5 external interface
  - F5 workload interface
  - F5 bastion interface
- Next steps

Setting up a web application firewall with F5 BIG-IP

- Network path
- Before you begin
- Provision CIS
- Configure CIS TLS Mode
- Configure CIS Global Load Balancer
- Configure CIS Range Application
- Enable license modules on the BIG-IP
- Import client certificate and key to BIG-IP
  - Import client key
  - Import client certificate
- Import server-side intermediate certificates to BIG-IP
- Configure ciphers
  - Create a cipher policy

    - Create a cipher group
- Create SSL profiles in BIG-IP
    - Create a client SSL profile
    - Create a server SSL profile
- Setting Up the HPCS integration with F5
- Create a BIG-IP virtual server
    - Create iRules
    - Create the backend pool
    - Create policies
    - Create the virtual server
- Configure WAF on the virtual server
    - Enable license
    - Set up WAF

Setting up full tunnel client-to-site VPN that uses F5 BIG-IP

- Objectives
    - Before you begin
- Enable license modules on the BIG-IP
- Set up an AAA Server in BIG-IP for VPN users
- Set up a client SSL profile
- Create a BIG-IP VPN
- Configure the VPN virtual server
- Test the BIG-IP VPN

Bastion host

Setting up a bastion host that uses Teleport

- Before you begin
- Limitations
- Configure Object Storage
- Configure App ID
- Provision virtual server instances for Teleport Enterprise
- Configure a security group for your virtual server instances
- Configure an access control list
- Configure virtual server instances for Teleport Enterprise
- Log in to the bastion host
    - Log in through the the web console
    - Log in through the tsh client
- Install tools
- Remove SSH port from security group and access control list
- Related controls in IBM Cloud Framework for Financial Services

General deployment and configuration information

Set up environment for deployment and configuration

- IBM Cloud Command Line Interface
- Application programming interfaces
- Terraform for IBM Cloud
- Next steps

Use deployable architectures to deploy your resources

- Next steps

Intra Organization Multitenancy Architecture Reference

- Next steps

Accounts, identity management, and access control

IBM Cloud account setup

- Next steps

Organizing IBM Cloud accounts and resources

- Organization for a single deployment within an account
  - VPC reference architecture
  - Satellite reference architecture
- Organization for multiple deployments
  - VPC reference architecture
  - Satellite reference architecture
- Scaling the number of deployments within an account
  - VPC reference architecture
  - Satellite reference architecture
- Managing scalability with an enterprise
  - Rationale for using an enterprise
  - Organizing your enterprise
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Access management in IBM Cloud

- Managing access with access groups
  - Example access group structure
- Invite users to your account
- Access groups with dynamic rules
- Considerations for enterprises
- Required permissions for services in reference architecture
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Handling and securing secrets

- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Consumer accounts for application provider workloads

- App ID integrations with IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Tagging IBM Cloud resources and managing access with tags

- Example tagging scheme
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Networking

Networking for VPC reference architecture

VPC networking overview

- Next steps

Creating and connecting the management and workload VPCs

- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Connecting application provider to the management VPC

- Direct Link
- VPN for VPC
- Full tunnel client-to-site VPN
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Running operator actions through a bastion host

- Set up a bastion host
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Consumer connectivity to workload VPC

- Consumer in same organization as application provider
  - Direct Link
  - VPN for VPC
- Consumer in different organization than application provider
  - Connecting from private intranet
  - Connecting from public internet
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

- Use cases
- Reference deployment architecture
- Application user access (non-privileged)
    - Security enforcements
    - Threat mitigation
    - Architectural decisions
- Cloud administrator access (privileged)
    - Security enforcements
    - Threat mitigation
    - Architectural decisions
- Application administrator access (privileged)
    - Security enforcements
    - Threat mitigation
    - Architectural decisions
- Cloud console access
    - Security enforcements
    - Threat mitigation
    - Architectural decisions
- Transit Gateway connectivity (same account)
    - Security enforcements
    - Threat mitigation
- Transit Gateway connectivity (cross-account)
    - Security enforcements
    - Architectural decisions
    - Threat mitigation
- Public Gateway access
    - Security enforcements
    - Threat mitigation
    - Architectural decisions
- Client-to-site VPN access
    - Security enforcements
    - Threat mitigation
    - Architectural decisions
- Site-to-site VPN access (ISV to cloud)
    - Security enforcements
    - Threat mitigation
    - Architectural decisions
- Site-to-site VPN access (service consumer to cloud)

Networking for Satellite reference architecture

Consumer connectivity to workload resources

- Edge plane with web application firewall
- Exposing your Red Hat OpenShift on IBM Cloud workloads
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Compute and containers

Working with virtual servers in VPC reference architecture

- Deployment
- Next steps

Working with Red Hat OpenShift on IBM Cloud

- Deploying Red Hat OpenShift on IBM Cloud
- Related resources
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Container image management

Image management in Red Hat OpenShift on IBM Cloud clusters

- Scope
  - Repository access
  - Image acceptance policies
  - Vulnerability scanning
  - Out of scope
- Approaches and recommendations
  - Explicit references to images in Container Registry private namespaces
  - Image mirroring to Container Registry for transparent image reference resolution
  - Operator index and bundle image mirroring to Container Registry for operator deployments
- Next steps

Guidance for Red Hat OpenShift on IBM Cloud administrators

- Outbound traffic protection and external registries
- Container image integrity
- Operator deployment
- Image mirroring
  - Option 1: Code Engine job to run the mirroring process
  - Option 2: Mirroring process on a VSI in Management VPC
  - Option 3: Mirroring process running within an enterprise environment outside of IBM Cloud
- Image and operator update strategy
  - Operator updates
  - Image updates

- Next steps

Configuring registry mirroring on Red Hat OpenShift on IBM Cloud nodes

- Understanding registry mirroring
- Before you begin
- Creating a DaemonSet for registry configuration update
    - Creating the mapping file
    - Creating the ConfigMap
    - Deploying the DaemonSet
- Creating a DaemonSet for updating node level Container Registry pull secret
    - Creating the pull secret
    - Deploying the DaemonSet
- Updating DaemonSets
- Next steps

Configuring operator mirroring and catalog sources

- Understanding operator mirroring
- Creating CatalogSource resources
- Applying the CatalogSource
- Verifying operator availability
- Next steps

Configuring image acceptance policies

- Installing Portieris
- Understanding policy types
- Creating a default deny-all cluster policy
- Creating namespace-specific image policies
- Verifying image signatures
- Example: Allowing Sysdig agent for Security and Compliance Center Workload Protection
- Next steps

Scanning container images for vulnerabilities in IBM Cloud

- SCC WP registry scanning
- SCC WP agent for scanning of images deployed on a cluster
- Container Registry scanning
- Next steps

Container image management references

- IBM Cloud documentation
- Red Hat Red Hat OpenShift on IBM Cloud documentation
- IBM Cloud for Financial Services
- Additional resources

Storage

Storage for VPC reference architecture

- IBM® Cloud Block Storage for Virtual Private Cloud
- IBM Cloud® Object Storage
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Storage for Satellite reference architecture

- Storage in IBM Cloud
    - About Object Storage
- Storage in Satellite location
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Logging and monitoring

Audit logging of IBM Cloud events

- Activity Tracker Event Routing
- Activity Tracker hosted event search
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Audit logging of application provider events and SIEM

- Auditable events
- Storing events in a SIEM
- Related IBM Cloud for Financial Services controls
- Next steps

Operational logging

- Application logging
- Platform logging
    - VPC network traffic
    - Red Hat OpenShift on IBM Cloud and virtual server instances
    - Other IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

IBM Cloud Logs best practices for Financial Services

- Understanding data retention
    - Recommended retention period
- Configuring IBM Cloud Logs storage
    - Creating storage buckets
    - Connecting buckets to IBM Cloud Logs
- Understanding data tiers
    - Data pipeline flow

- Sending logs to IBM Cloud Logs
  - Activity Tracker event routing (required)
  - Application logs (required)
  - IBM Cloud service logs (required)
  - IBM Cloud Logs extensions (optional)
- Using the TCO Optimizer
  - Recommended TCO policies
  - Accessing the TCO Optimizer
  - Creating a TCO policy
- Next steps
- Related information

Compliance monitoring

- Using Compliance Manager
- Compliance best practices for Financial Services Validated services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Operational monitoring

- Red Hat OpenShift on IBM Cloud
- Virtual Servers for VPC
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Data encryption

Data encryption at rest

- Data encryption at rest in IBM Cloud
  - Setting up Hyper Protect Crypto Services
  - Setting up Hyper Protect Crypto Services with Bring Your Own HSM
  - Setting up Key Protect
  - Data encryption at rest in Satellite location
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Data encryption in transit

- TLS requirements
  - TLS protocol requirements
  - Key size requirements
  - Acceptable cipher suites
  - Certificate requirements
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Business continuity and disaster recovery

Business continuity and disaster recovery overview

- Requirements
    - Alternative storage site
    - Information system backup
    - Information system recovery and reconstitution
- Backup and disaster recovery for the reference architectures
- Backup and disaster recovery for IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Business continuity and disaster recovery for VPC reference architecture

- Workloads in VPC reference architecture
    - Workloads on virtual server instances
    - Workloads on Red Hat OpenShift on IBM Cloud
- Backup and disaster recovery for IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Business continuity and disaster recovery for Satellite reference architecture

- BCDR in your Satellite location
- Backup and disaster recovery for IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

High availability

High availability overview

- HA for the reference architectures
- HA in IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

High availability for VPC reference architecture

- Your workloads on VPC infrastructure
    - Workloads in multizone regions
    - Workloads in multiple regions
- HA for IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

High availability for Satellite reference architecture

- High availability in your Satellite location

    - Your workloads in Red Hat OpenShift on IBM Cloud
- HA for IBM Cloud services
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Development processes and software integrity

Development processes and software integrity

- Red Hat OpenShift on IBM Cloud
- Virtual server instances
- Related controls in IBM Cloud Framework for Financial Services
- Next steps

Responsibilities for operating services in your deployment

Responsibilities for operating services in your deployment

- Responsibilities for Financial Services Validated services

IBM Cloud for Financial Services Validation Process for Partners

IBM Cloud for Financial Services Validation Process for Partners

IBM Cloud for Financial Services Validated Partner Offerings

IBM Cloud for Financial Services Validated Partner Offerings

- List of FS Validated ISV Offerings

FAQs for IBM Cloud for Financial Services

FAQs for IBM Cloud for Financial Services

- What are the supported reference architectures for the IBM Cloud for Financial Services?
- What is a *consumer* and what is an *application provider*?
- What does it mean for a service to be IBM Cloud for Financial Services Validated?
- Should I only use IBM or third-party managed services that are IBM Cloud for Financial Services Validated?
- What if I don't want to use managed services, but I want to install and manage my own software?

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
2026-07-30