

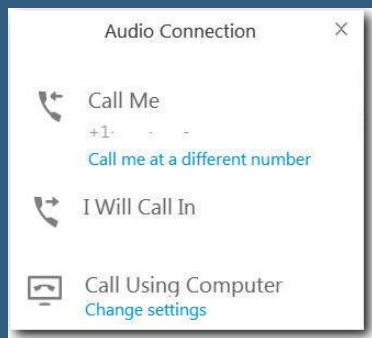
# QRadar: Cloud Architecture

Capabilities, collection, and best practices panel

## IBM SECURITY SUPPORT OPEN MIC

To hear the WebEx audio, **select an option** in the Audio Connection dialog or by access the Communicate > Audio Connection menu option. To ask a question by voice, you must either Call In or have a microphone on your device. *You will not hear sound until the host opens the audio line.*

For more information, visit:  
[http://ibm.biz/WebExOverview\\_SupportOpenMic](http://ibm.biz/WebExOverview_SupportOpenMic)



**NOTICE:** BY PARTICIPATING IN THIS CALL, YOU GIVE YOUR IRREVOCABLE CONSENT TO IBM TO RECORD ANY STATEMENTS THAT YOU MAY MAKE DURING THE CALL, AS WELL AS TO IBM'S USE OF SUCH RECORDING IN ANY AND ALL MEDIA, INCLUDING FOR VIDEO POSTINGS ON YOUTUBE. IF YOU OBJECT, PLEASE DO NOT CONNECT TO THIS CALL.

# Disclaimer

## Please Note:

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.



# Third Party Cloud Vendors

# Cloud Installs

## Currently Supported



## Planned



Google Cloud Platform



# Cloud Ingestion

## Currently Supported



## Planned



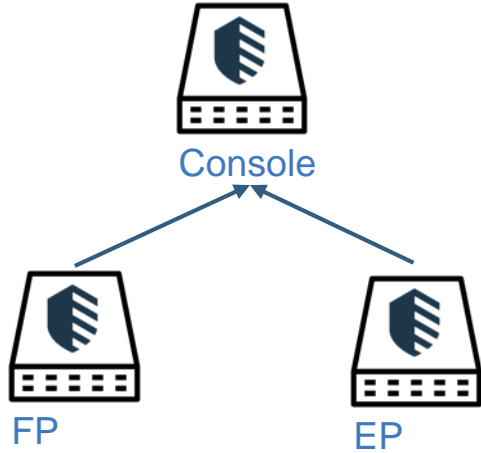
Google Cloud Platform



# AWS Deployment Architecture Examples


# Collect from the Cloud – AWS Infrastructure Logging

QRadar On Premise



EC = Event Collector  
EP = Event Processor  
FC = Flow Collector  
FP = Flow Processor  
← - REST API

AWS Ingestion Example



REST API

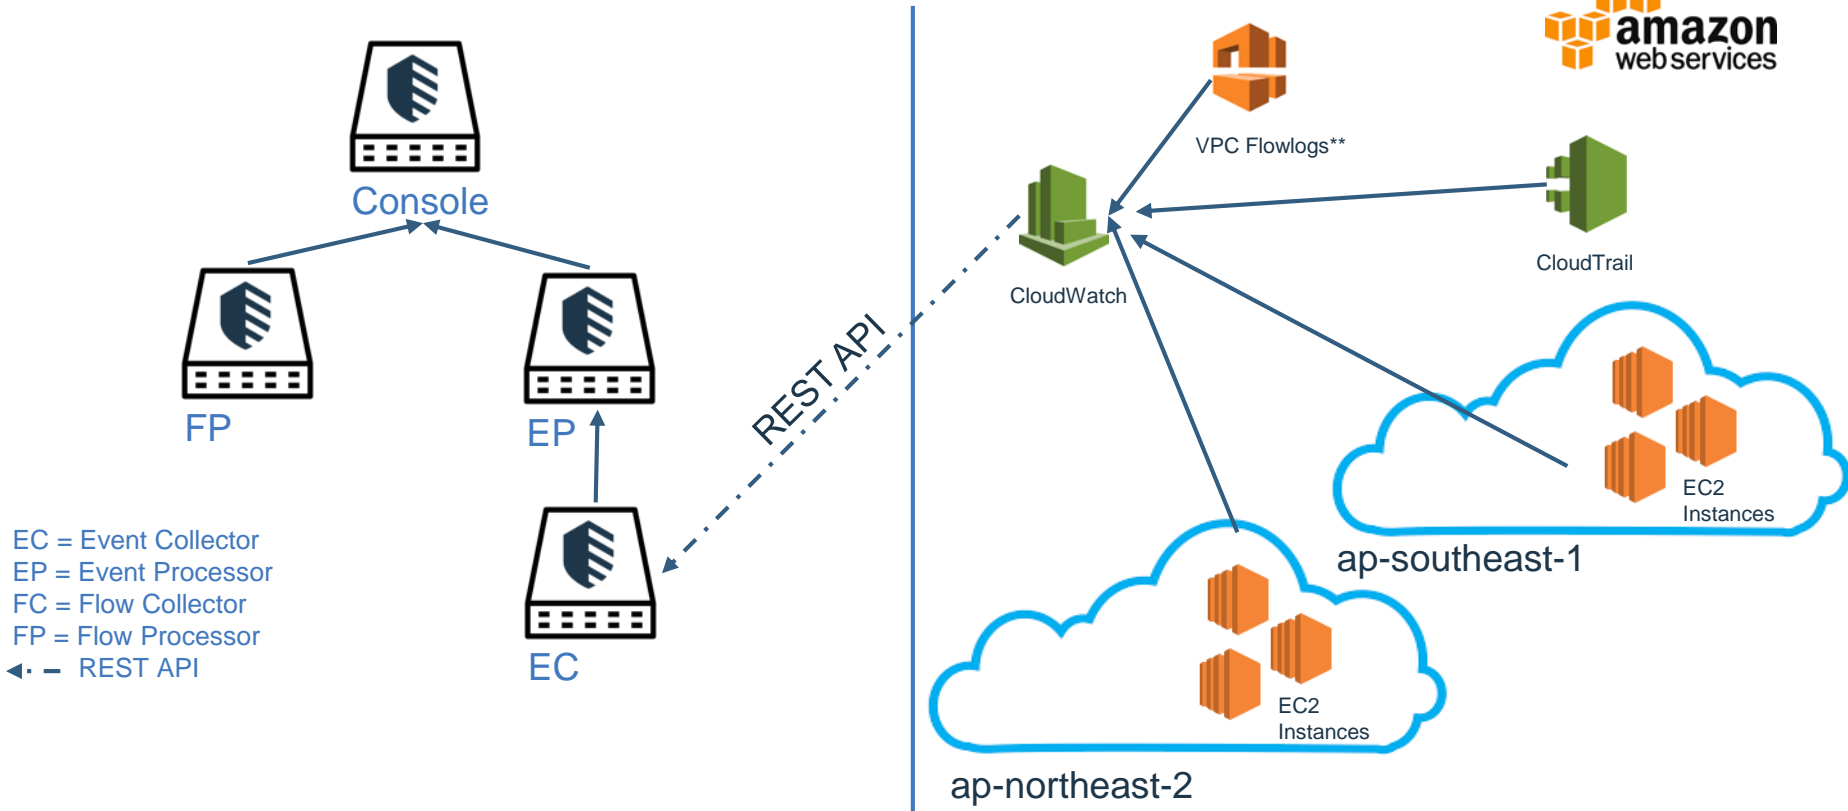
REST API

This diagram outlines how event collection works for most users today. Data is sent to the QRadar On Premise appliances. \*\*VPC Flowlogs are captured as events now.

# Collect from the Cloud – AWS Collection Example 1

QRadar On Premise

AWS Multi-Region Example



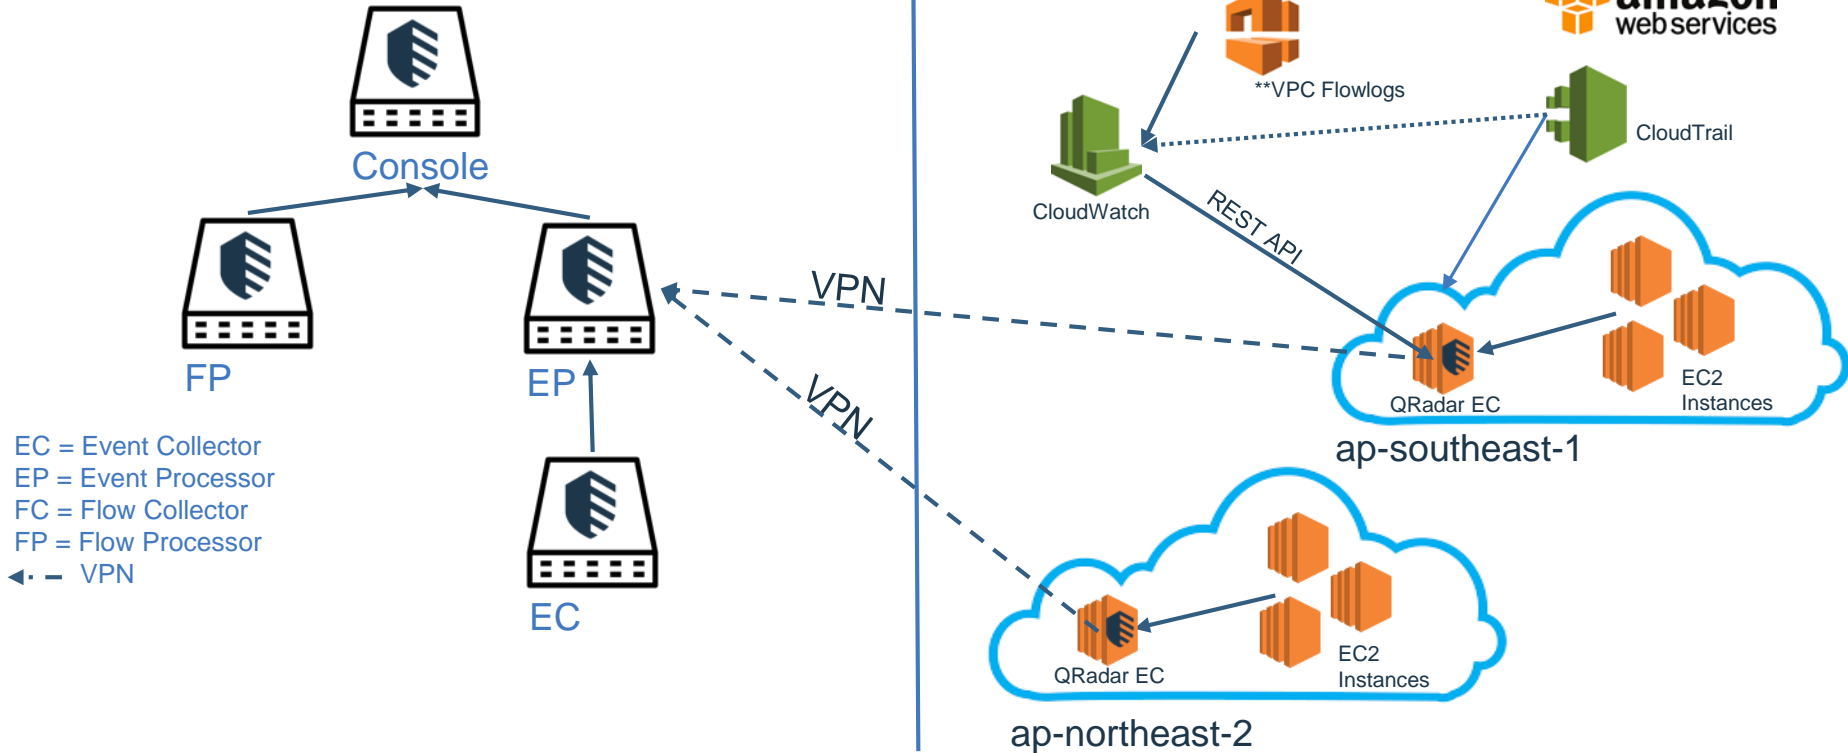
In this architecture outline data from ECS instances and CloudTrail events are sent to CloudWatch for collection. The QRadar Amazon REST API collects the CloudWatch events and VPC Flowlog events\*\*.



# Collect from the Cloud – AWS Collection Example 2

QRadar On Premise

AWS Multi-Region Example






This diagram outlines collection with an Event Collector in the Cloud. Data is collected on the Cloud EC and the QRadar pipeline sends the data to an EP On Premise appliance. A benefit is that you save on bandwidth as the event pipeline compresses EC to EP connections. Searches are completed with the on premise appliances.

# Collect from On Premise and Forward to QRadar Cloud



QRadar On Premise

AWS Multi-Region Example



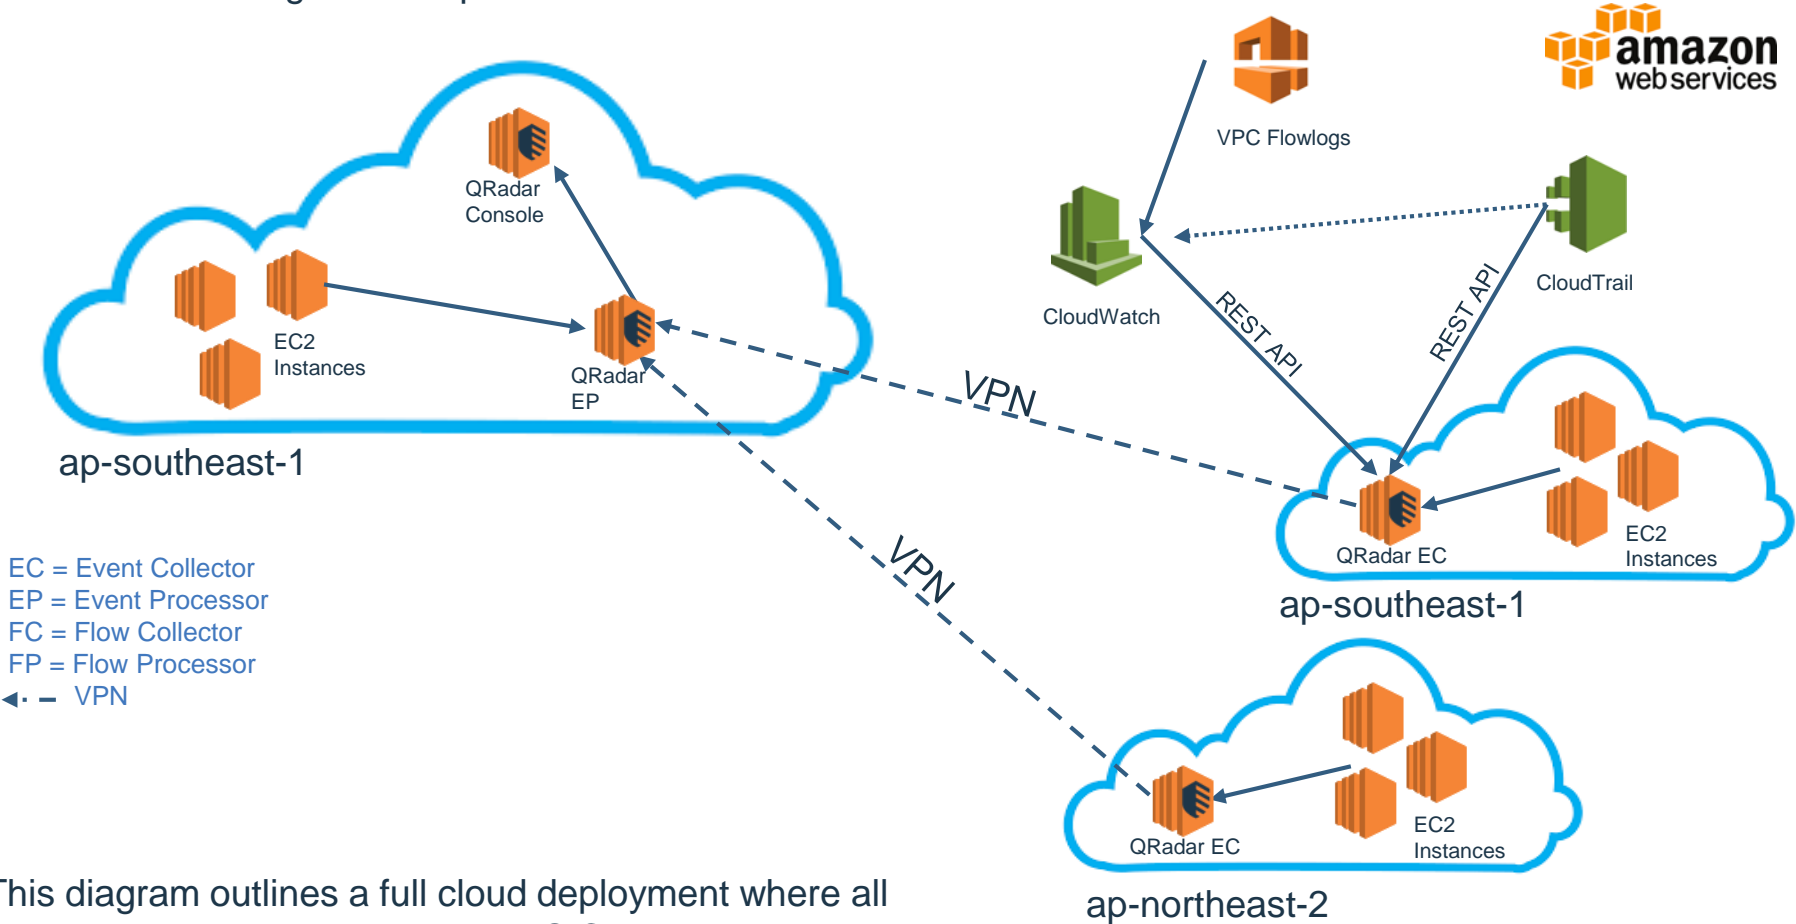
EC = Event Collector  
EP = Event Processor  
FC = Flow Collector  
FP = Flow Processor  
← - VPN

This example outlines on premise Event Collector and Flow Collector appliances that VPN data to the QRadar Cloud deployment. A benefit here is it limits the requirement for HA due to Cloud resiliency.



# Collect from the Cloud – Full Cloud Deployment in AWS

## AWS Multi-Region Example



This diagram outlines a full cloud deployment where all appliances are installed in the AWS Cloud.

# Frequently Asked Questions for AWS Installations

## **Q1: Is QRadar HA supported in AWS?**

A1: Not at the moment. The resiliency is provided by the cloud vendor.

## **Q2: In a hybrid deployment should I deploy an EC or an EP in AWS?**

A2: It is recommended to keep the EP in the same area as the console. This helps search performance and it makes the data egress charges from AWS deterministic. You will send all data out of AWS at an approximate 10:1 compression ratio.

## **Q3: Do I have to deploy an EC or Data Gateway in AWS to collect logs from AWS?**

A3: No. CloudTrail or CloudWatch Logs can be collected from anywhere. It is possible to send EC2 instance logs (OS and application) to CloudWatch Logs.

## **Q4. What considerations are required for my network hierarchy?**

A4. Administrators need to decide if they consider the hosts generating events in AWS your assets or not. Depending on the applications or OS and how they are used might make a difference.

## **Q5. How do I address domain separation & overlapping IPs for multiple cloud environments?**

A5. Administrators need to think about how they implement domains as they move assets or infrastructure in the cloud and how IP addresses will overlap.

# Frequently Asked Questions for AWS Event Collection

**Q6: Can QRadar collect logs in AWS Cloudtrail from the root directory?**

**For example, /AWSLogs instead of /AWSLogs/<AccountNumber>/CloudTrail/<Region>/?**

A6: No, we can not use the root directory because we need to be able to identify the accounts.

**Q7: Will AWS Role Based Access be supported?**

A7: Yes, administrators or users interested in role based access can talk to use about a beta of this feature that is available.

**Q8. What are Virtual Private Cloud (VPC) Flows?**

A8. VPC Flows capture information about the IP traffic going to and from network interfaces in your VPC.

**Q9. What VPC Flow logs events supported?**

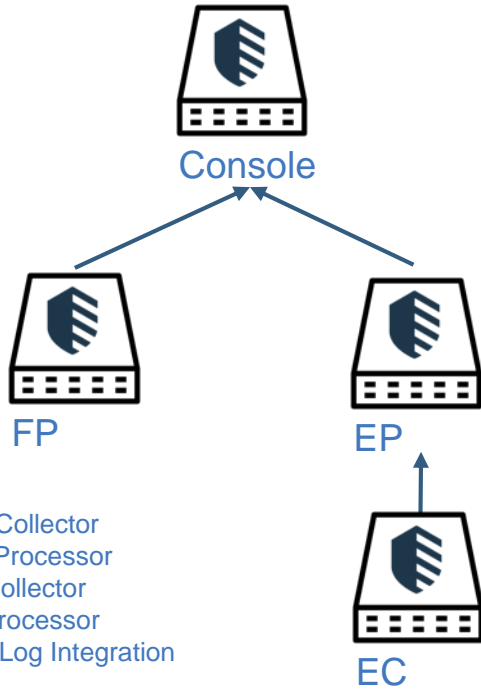
A9. The traffic going in and out of your network interfaces in your Amazon VPC. The traffic actions basically ACCEPT and REJECT



# Azure Deployment Architecture Examples

# Collect from the Cloud – Azure Infrastructure and Instance Logging

QRadar On Premise



EC = Event Collector  
EP = Event Processor  
FC = Flow Collector  
FP = Flow Processor  
ALI = Azure Log Integration  
← - TLS

Azure Ingestion Example



Links:

- [Getting started with Azure log integration](#)
- [QRadar DSM Guide Azure log integration](#)

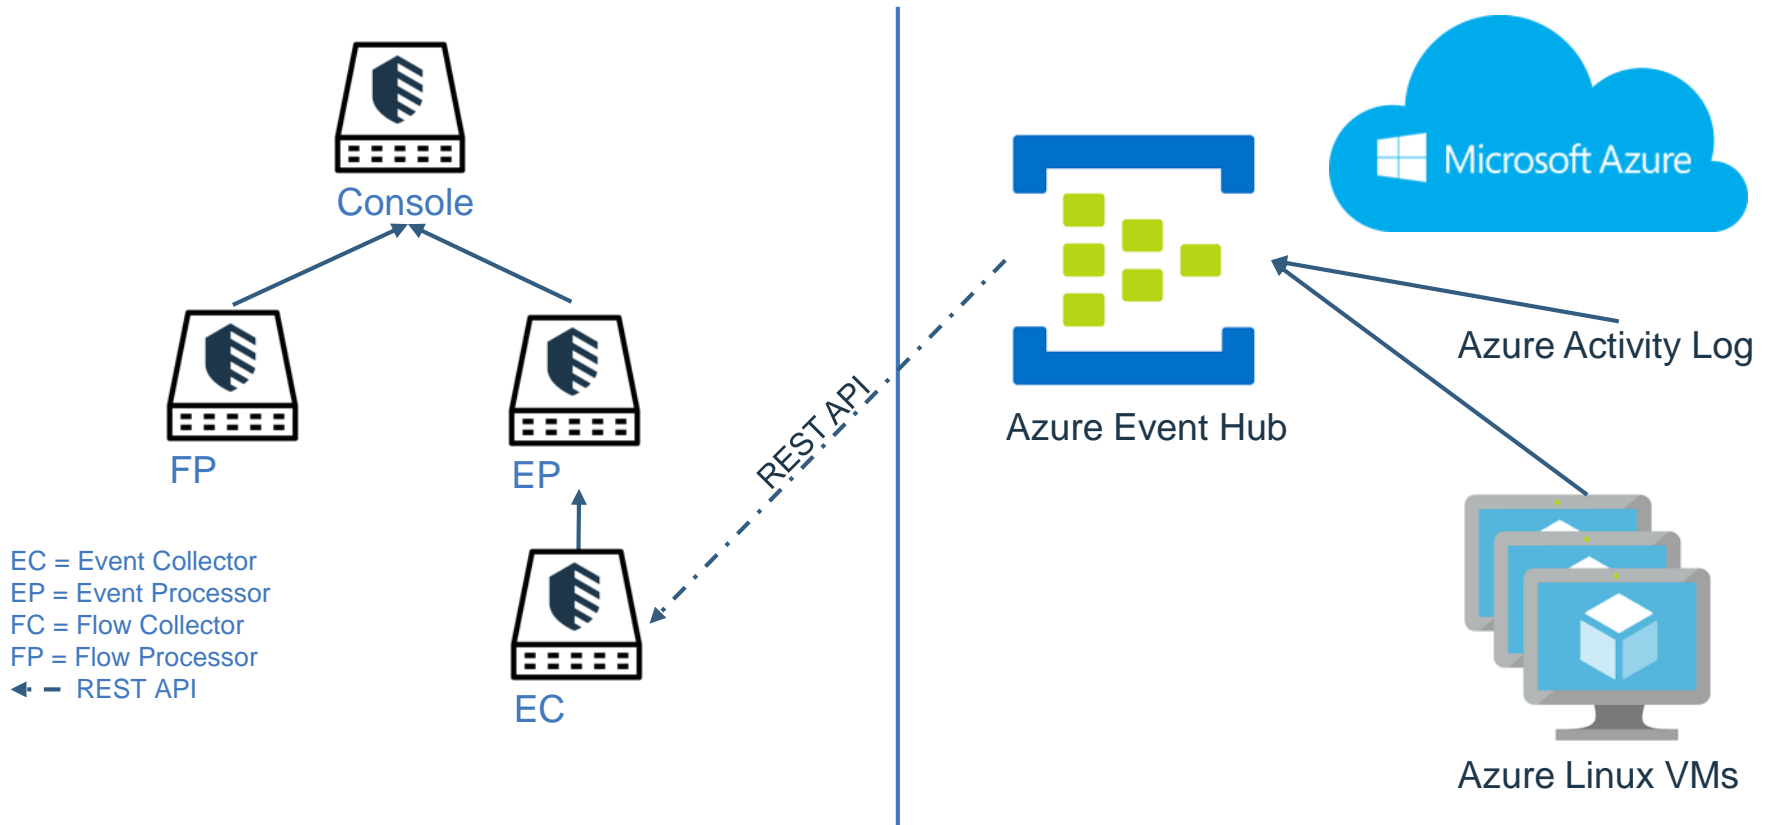
This diagram outlines how event collection works for Azure Activity Logs today. Data is sent to the Azure Log Integrator and queried by the QRadar On Premise appliance using a REST API.

\*\* The Azure log integration service can be on premise or in the Azure Cloud (See Getting Started)

# Collect from the Cloud – Azure Infra and Instance Logging

QRadar On Premise

Azure Ingestion Example

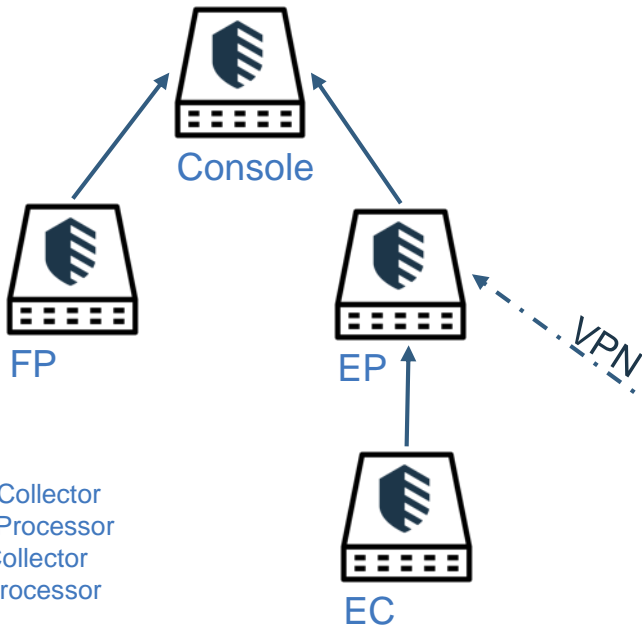


This diagram outlines how event collection works for the Azure Event Hub. Data is sent to the Azure Event Hub and queried by the QRadar On Premise appliance using a REST API.



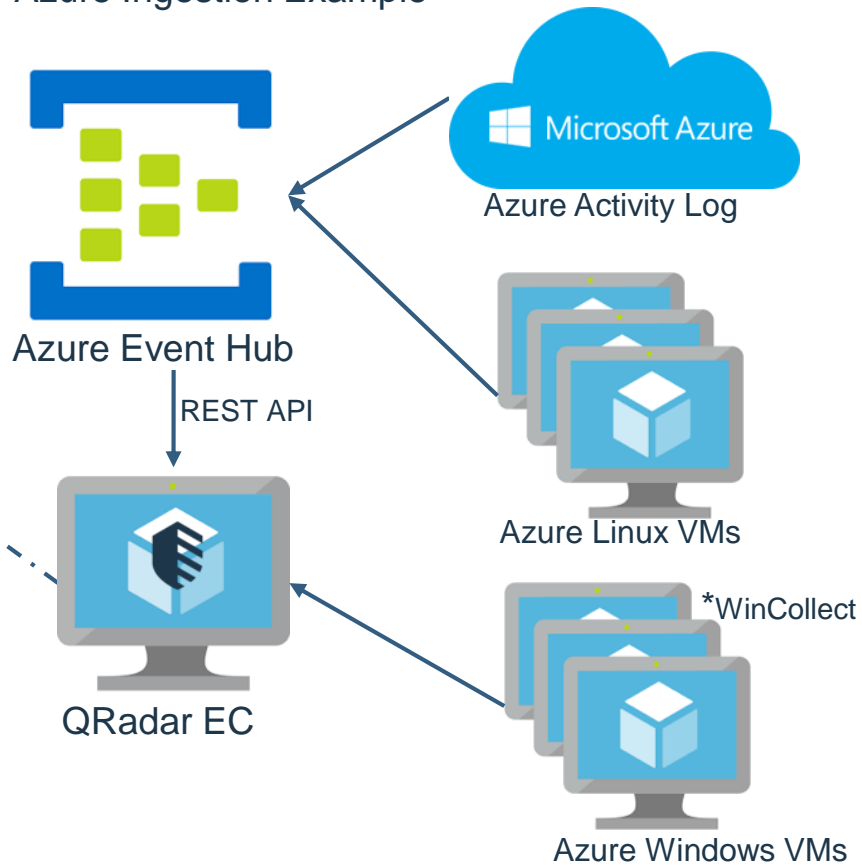
# Collect from the Cloud – Azure Infra and Instance Logging (Future)

QRadar On Premise



EC = Event Collector  
EP = Event Processor  
FC = Flow Collector  
FP = Flow Processor  
← - VPN

Azure Ingestion Example



This diagram outlines an architecture concept for hosting a QRadar Event Collector in the Azure Cloud.

# Frequently Asked Questions for Azure Installations

## **Q1: Can I install QRadar in Azure today?**

A1: Not at the moment.

## **Q2: Why is QRadar not available in Azure?**

A2: QRadar requires the base version of RHEL with no package changes to install on. This is not available in the Azure marketplace.

## **Q3: Will QRadar HA be supported in Azure?**

A3: Not at the moment. The resiliency is provided by the cloud vendor.

## **Q4: Does the Azure Event Hub Protocol support Windows events?**

A4: No, the Azure Event Hub Protocol does not support Windows events. The solution at the moment is to use WinCollect agents.

# Frequently Asked Questions for Azure Event Collection

## **Q5. Is proxy supported?**

A5. No, Proxy settings are not implemented in the design of this protocol because AMQP uses TCP to connect and the Microsoft Azure Event Hubs Event Processor does not provide any method to connect via proxy to bypass this.

## **Q6. Does the Azure Event Hub Protocol support Windows events?**

A6. No, the Azure Event Hub Protocol does not support Windows events. The solution at the moment, is to use WinCollect agents.

## **Q7. What kind of events can the protocol handle?**

A7. Azure Event Hub collects data in the following categories; Azure Activity Logs, Diagnostic Logs, Linux Events and generic Syslog events. Azure Activity and Diagnostic logs are received as JSON and are very similar to each other, both use the same payload format. Both of these event types are handled by the Microsoft Azure DSM. The Linux Events are JSON formatted and converted to syslog so that auto discovery can figure out which Linux event type it falls under (DHCP server, iptables firewall or OS). Generic syslog events are received as per syslog format.

## **Q. What is the retention period to store events?**

A8. Azure Event Hubs can collect events and then store them for a user configurable retention period, the current maximum retention period is 7 days.



# Installing QRadar in AWS Today

# Installing QRadar in AWS - Today

1. Choose your image: RHEL-7.3\_HVM\_GA-20161026-x86\_64-1-Hourly2-GP2 from Community AMIs
2. Choose your EC2 instance (M4.2XLarge or above based on [Virtual Appliance Sizing Guide](#))
3. Choose 100GB for the root disk (GP2 is fine)
4. Choose an appropriate size for the secondary disk(s) based on EPS average Payload Size and Retention
  - Disks can be either GP2 or IO1 disks. IO1 with the appropriately provisioned IOPs is recommended.
  - LVM is supported now, so you can start small and expand storage as needed by adding more disks
  - Optionally you can later expand storage using Data Nodes and new EC2 instances to scale storage and search speed
5. Setup your security group to allow port 22 and 443 to a set of whitelisted IPs
6. Choose your key pair or create one
7. Review and Launch the Instance
8. As the ec2-user scp over the aws\_qradar\_prep.sh script and the ISO  
  
Example: `scp -i <key.pem> aws_qradar_prep.sh ec2-user@<public ip>:`
9. As root run `aws_qradar_prep.sh --install`, then mount the ISO and run `/media/cdrom/setup`
10. Use the internal IPs for the network configuration
11. Estimated 1-2hrs from start to finish

# Automating Some Installation Steps with User Data

## QRadar

- Create an S3 bucket and upload the QRadar ISO and aws\_qradar\_prep.sh script
- Create an IAM Role with S3 Read Only permissions
- When launching the EC2 instance give the Instance the IAM role and enter the following in User Data:

```
#!/bin/bash
# Install the awscli and get the ISO from your S3 bucket
yum install -y python-setuptools
easy_install awscli
aws s3 cp s3://<s3bucket>/Rhe764QRadar7_3_1_20171206222136.stable-7-3-1.iso /home/ec2-user/qradar.iso
aws s3 cp s3://<s3bucket>/aws_qradar_prep.sh /home/ec2-user/

# Update dracut (for QRadar 7.3.1) and run the prep script
yum update -y dracut
mkdir /media/cdrom
bash +x /home/ec2-user/aws_qradar_prep.sh --install
```



# Installing QRadar CE in AWS

# Installing QRadar Community Edition - Today

1. Choose the Centos 7 image from the AWS Marketplace
2. Choose your EC2 instance (T2.Medium or above according to the [Community Edition Install Guide](#))
3. Choose 100GB for the root disk or larger (no real need for a secondary disk unless you want to separate data store from the instance root volume)
4. Setup your security group to allow port 22 and 443 to a set of whitelisted IPs
5. Choose your key pair or create one
6. Review and Launch the Instance
7. As the centos user scp over the ISO:  
Example: `scp -i <key.pem> QRadarCE7_3_0_20171013140512.GA.iso centos@<public ip>:`
8. As root mount the ISO and run `/media/cdrom/setup`
9. Use the internal IPs for the network configuration
10. Estimated 1-2hrs from start to finish



# Automating Some Installation Steps with User Data

## QRadar

- Create an S3 bucket and upload the QRadar CE ISO
- Create an IAM Role with S3 Read Only permissions
- When launching the EC2 instance give the Instance the IAM role and enter the following in User Data:

```
#!/bin/bash
# Install the awscli and get the ISO from your S3 bucket
yum install -y python-setuptools
easy_install awscli
aws s3 cp s3://<s3bucket>/QRadarCE7_3_0_20171013140512.GA.iso /home/centos/qradar.iso

# Make the cdrom dir and mount the iso
mkdir /media/cdrom
mount -o loop /home/centos/qradar.iso /media/cdrom
```



# Installing QRadar in AWS (Soon)

# Installing QRadar in AWS - Soon

1. Choose the QRadar Console AMI or QRadar Managed Host AMI from the AWS Marketplace
2. Choose your EC2 instance (M4.2XLarge or above based on [Virtual Appliance Sizing Guide](#))
3. If it's a managed host enter the type of managed host in User Data
4. Choose 100GB for the root disk (GP2 is fine)
5. Choose an appropriate size for the secondary disk(s) based on EPS average Payload Size and Retention
  - Disks can be either GP2 or IO1 disks. IO1 with the appropriately provisioned IOPs is recommended.
  - LVM is supported now, so you can start small and expand storage as needed by adding more disks
  - Optionally expand storage using Data Nodes and new EC2 instances to scale storage and search speed
6. Setup your security group to allow port 22 and 443 to a set of whitelisted IPs
7. Choose your key pair or create one
8. Review and Launch
9. Estimated **10-15 minutes** from start to finish



# Instance Log Ingestion from Auto-Scaling Groups

# 65,534 problems

## Log Source Admin

- Default VPC size is a /16 in AWS, that's 65,534 useable IPs
- EC2 instances sending logs to QRadar could live for minutes, days, months, or years
- Over time with an auto-scale group you could create 65,534 log sources (identified by internal IP) of which the majority are going to be inactive
- Autodetection may be difficult for some Linux OS sources and manually creating the log source per IP is not feasible

## Uniqueness

- Your internal IP is not unique and may be re-used over time, perhaps within the same day by a separate instance which may have a different application or OS
- The OS logs in an EC2 instance have only the internal IP context and knows nothing about the cloud it is running in
- The cloud meta-data is really what defines a unique instance (instance id, interface id, account, et cetera)

# RSyslog Solution For Linux Instances

- Use one log source identifier for an auto-scale group or application
- Create an Rsyslog Template to alter the hostname in the header to match the log source identifier of your choice
- Insert the cloud meta data between the syslog header and the payload
- Automate all of this with User Data on EC2 Instance Launch

rsyslog template

```
template(name="RFC3164ForwardFormat" type="list") {
  constant(value("<")
  property(name="pri")
  constant(value(">")
  property(name="timestamp")
  constant(value=" ")
  constant(value="LinuxAppAlpha")
  constant(value=" ")
  constant(value="instanceId: INSTANCEID, ")
  constant(value="accountId: ACCOUNTID, ")
  constant(value="interfaceId: INTERFACEID, ")
  property(name="syslogtag" position.from="1" position.to="32")
  property(name="msg" spifnolstsp="on" )
  property(name="msg")
}
```

```
$ActionForwardDefaultTemplate RFC3164ForwardFormat
authpriv.* @@QRADARIP:514
```

# RSyslog Solution For Linux Instances - continued

## userdata script

```
#!/bin/bash
export PATH=~/.local/bin:$PATH
curl -O https://bootstrap.pypa.io/get-pip.py
python get-pip.py --userpip install awscli --upgrade -user

TEMPLATENAME=qradarforwardingtemplate.conf
TEMPLATEFILE=/etc/rsyslog.d/$TEMPLATENAME

INSTANCEID=$(curl http://169.254.169.254/latest/meta-data/instance-id 2>/dev/null)
ACCOUNTID=$(curl http://169.254.169.254/latest/dynamic/instance-identity/document 2>/dev/null |
python -c 'import sys, json; print json.load(sys.stdin)["accountId"]')
MAC=$(curl http://169.254.169.254/latest/meta-data/mac 2>/dev/null)
INTERFACEID=$(curl http://169.254.169.254/latest/meta-data/network/interfaces/macs/$MAC/interface-id
2>/dev/null)

aws s3 cp s3://<s3bucket>/$TEMPLATENAME $TEMPLATEFILE

sed -i s/INSTANCEID/$INSTANCEID/ $TEMPLATEFILE
sed -i s/ACCOUNTID/$ACCOUNTID/ $TEMPLATEFILE
sed -i s/INTERFACEID/$INTERFACEID/ $TEMPLATEFILE

sed -I s/QRADARIP/<qradarip>/ $TEMPLATEFILE

service rsyslog restart
```



# Resources



# Resources

## QRadar on Cloud

- <https://www.ibm.com/us-en/marketplace/hosted-security-intelligence>
- [https://www.ibm.com/support/knowledgecenter/en/SSKMKU/com.ibm.QRadar.doc\\_cloud/c\\_QRadar\\_hosted\\_overview.html](https://www.ibm.com/support/knowledgecenter/en/SSKMKU/com.ibm.QRadar.doc_cloud/c_QRadar_hosted_overview.html)

## QRadar and AWS

- [https://www.ibm.com/support/knowledgecenter/en/SS42VS\\_7.3.0/com.ibm.QRadar.doc/t\\_Cloud\\_Install\\_QRadar\\_AWS.html](https://www.ibm.com/support/knowledgecenter/en/SS42VS_7.3.0/com.ibm.QRadar.doc/t_Cloud_Install_QRadar_AWS.html)
- [https://www.ibm.com/support/knowledgecenter/en/SS42VS\\_DSM/c\\_dsm\\_guide\\_amazon\\_aws\\_ct\\_overview.html](https://www.ibm.com/support/knowledgecenter/en/SS42VS_DSM/c_dsm_guide_amazon_aws_ct_overview.html)
- <https://exchange.xforce.ibmcloud.com/hub/extension/bf358419d91d425df1e2ee9e72d37c13>

## QRadar and Azure

- <https://blogs.msdn.microsoft.com/azuresecurity/2016/09/24/integrate-azure-logs-to-QRadar/>
- [Getting started with Azure log integration \(https://docs.microsoft.com/en-us/azure/security/security-azure-log-integration-get-started\)](https://docs.microsoft.com/en-us/azure/security/security-azure-log-integration-get-started)
- [QRadar DSM Guide Azure log integration](#)

## OpenVPN Configuration

- [https://www.ibm.com/support/knowledgecenter/en/SS42VS\\_7.3.0/com.ibm.QRadar.doc/t\\_cloud\\_server\\_vpn\\_.html](https://www.ibm.com/support/knowledgecenter/en/SS42VS_7.3.0/com.ibm.QRadar.doc/t_cloud_server_vpn_.html)
- [https://www.ibm.com/support/knowledgecenter/en/SS42VS\\_7.3.0/com.ibm.QRadar.doc/t\\_cloud\\_client\\_vpn.html](https://www.ibm.com/support/knowledgecenter/en/SS42VS_7.3.0/com.ibm.QRadar.doc/t_cloud_client_vpn.html)
- [https://www.ibm.com/support/knowledgecenter/en/SS42VS\\_7.3.0/com.ibm.QRadar.doc/t\\_cloud\\_member\\_vpn.html](https://www.ibm.com/support/knowledgecenter/en/SS42VS_7.3.0/com.ibm.QRadar.doc/t_cloud_member_vpn.html)




# THANK YOU

## FOLLOW US ON:

 [facebook.com/IBMSecuritySupport](https://facebook.com/IBMSecuritySupport)

 [youtube/user/IBMSecuritySupport](https://youtube/user/IBMSecuritySupport)

 [@askibmsecurity](https://twitter.com/askibmsecurity)

 [SecurityLearningAcademy.com](https://SecurityLearningAcademy.com)

 [securityintelligence.com](https://securityintelligence.com)

 [xforce.ibmcloud.com](https://xforce.ibmcloud.com)

© Copyright IBM Corporation 2018. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and other IBM products and services are trademarks of the International Business Machines Corporation, in the United States, other countries or both. Other company, product, or service names may be trademarks or service marks of others.

Statement of Good Security Practices: IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a lawful, comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM does not warrant that any systems, products or services are immune from, or will make your enterprise immune from, the malicious or illegal conduct of any party.