



Section 3: Qradar on Cloud (QRoC)

CERT PREP FOR TECHNICAL SALES FOUNDATIONS FOR IBM QRADAR FOR CLOUD (QROC) V1



What to watch for?

- Lots of content – don't drown in it.
- Look for the “Learning Point Star”



QRadar on Cloud – QRoC

- 
- WW Technical Sales Enablement
 - IBM Security

What does QRadar on Cloud (QroC) do?

IBM is using its QRadar Security Intelligence technology to the cloud in a bid to help companies prioritize major security threats more quickly and free up critical resources to fight cyber attacks.



This method allows companies to deploy QRadar on Cloud instead of an on premise solution.

- Improved Time to value
- Reduced Implementation and IT Management overhead




The Need for QRadar on Cloud - QRoC




Why SIEM in the Cloud?

Security information and event management delivered as a service




Why SIEM in the Cloud?

Security information and event management delivered as a service




Why SIEM in the Cloud?

Security information and event management delivered as a service



Why SIEM in the Cloud?

Security information and event management delivered as a service





Deploying QRadar on Cloud - QRoC



What is QRadar On Cloud? - Highlights

- QRadar as a Service served from IBM Cloud (the IaaS formerly known as IBM Cloud)
- Dedicated operations group managing infrastructure and QRadar components
 - System provisioning/upgrades
 - Availability monitoring
 - Backend administration activities (user provisioning/etc)
- HA/DR are standard
- Data is encrypted in flight and at rest
- Priced by EPS and retention (default retention is 90 days)



Comparing QRadar On-Premise and QRadar on Cloud

Service Component	On-Premises	QRadar on Cloud
Cap-Ex budget item	✓	
Op-Ex budget item	✓	✓
IBM installation, deployment and upgrade		✓
IBM professionally managed infrastructure		✓
System Health Monitoring		✓
Configure data collection (DSMs)	✓	✓
Compliance reporting	✓	✓
Advanced attack detection	✓	✓
Incident detection and management	✓	✓
Asset modeling and vulnerability correlation	✓	✓
QVM, QFlows	✓	✓
QRM	✓	
QNI	✓	

QRoC vs On-Premise QRadar (cont)

- QRadar On Cloud CAN scale
 - POC underway in excess of 100K EPS
 - Responded to deals in excess of 200K EPS and 3.2M FPM
- Full QRadar administration requires QRadar operations team interaction
 - User Management
 - Token generation
 - etc.
- QRadar On Cloud is always at the latest QRadar release




Where are we currently deployed



Montreal - Canada
Toronto - Canada
Dallas - USA
San Jose - USA
Sao Paulo - Brazil
London - UK
Frankfurt - Germany




Architectural View



- QRoC is offered as a Highly Resilient Solution served from IBM Cloud
- Offered as a single Tenant Solution by default on IBM Cloud Bare Metal Servers or VM depending on EPS
- Deployed in a virtual deployment if EPS is below 8K EPS
- Deployed on IBM Cloud Bare Metal Servers if above 8K EPS
- Retention requirements can be met with Data Nodes
- On premise data gateways can be deployed to provide a secure channel to transfer log events to the QRadar environment

QRoC – Automation & Supportability





QRoC – Compliance : Coverage Today

- What do we have today?
 - IBM Internal Security Standards
 - ITCS104/ITCS300/ITSS
- What does that mean?
 - Information Security Management System.
 - Best practices from IT security perspective.
- What are the focus areas?
 - ✓ Privileged User Management
 - ✓ Network & Infrastructure Security Reviews
 - ✓ Vulnerability scanning & monitoring
 - ✓ PSIRT adherence / regular patch monitoring
 - ✓ Penetration Testing
- How is it enforced/policed?
 - Monthly Self-Assessment by non-product group (security services)
 - Rolled into division wide score card

QRoC – Compliance : Coverage (Future)



- Risk Management Framework
 - Aligns with IBM standards
- Adoption underway (2018)


- EU Data requirement
 - May 2018
- QRoC onboard with IBM adoption plans
- Leverage learnings/approach for on premise customers?



QRadar on Cloud – QRoC Onboarding



QRoC – Onboarding Process & Timeline





QRadar on Cloud onboarding

- **Primary user of the system (admin)**

- Name:
IBM Web ID:

Additional users

Name:
IBM Web ID:

Data Gateways:

Number: (we need to create an auth token for each gateway you will add to your networks)

Internal IP(s) for each Data Gateway: (this is the IP address you will provision on your local network for the data gateway, if you are adding multiple gateways please provide all their IPs)

Time zone:

- System Time: (the best time zone for the console to be configured with):

IP Whitelist:

- Whitelist: (The IP range that your users and data gateways will be connecting from. This can be a list of individual addresses and/or CIDRs)



QRadar on Cloud onboarding (cont)

- After you purchase IBM® QRadar® on Cloud, IBM sends you the information required for you to use QRadar on Cloud.
- IBM will send you an email after you have purchased QRadar on Cloud. This email contains a link to the Gateway Landing Page.
- The following list describes information about Gateway:
 - Your QRadar on Cloud token. You need a token for each Gateway appliance that you want to use to connect to QRadar on Cloud on the IBM cloud.
 - A download link to the IBM Security QRadar ISO for your gateway appliance.
 - A copy of Red Hat Enterprise Linux (RHEL) only if your organization requires changes to the default partitions that the QRadar ISO configures when installed.
 - The software installation activation key for each gateway appliance.
 - The public Host Name of the console that you connect to through the gateway appliance.
 - The required licenses for your 6 QRadar on Cloud users.
 - Each gateway appliance in your deployment must have a unique Host Name.
- IBM provides you with two IP addresses for your QRadar on Cloud deployment. One is for the Console, and the second is for the VPN.
- Keep port 443 outbound open for these two IP addresses.



QRoC – Administration and Licensing



QRadar on Cloud - Administration

- Full Admin




- The customer does not get full 'admin' access.
- Only the DevOps group has full admin.
- The customer does not get command line access to the deployment in IBM Cloud

- SaaS Admin



- The SaaS Admin has reduced access to the Admin tab
- The role is added to the QRadar Deployment via a special SaaS RPM which does not ship with the on-premises product
- The customer can email sisaasop@ca.ibm.com for admin activity




SaaS Admin View

Service Levels to meet customer needs

- **A simple service level structure with flexible upgrade to meet the needs of a wide variety of customers**
- **Charge metric: EPS (Events per Second)**
 - EPS is the major charge metric currently used by on-premise QRadar product.
 - Consistency between on-premise and SaaS to avoid confusion and allow future migration.
- **Support multiple service levels**
 - Basic Service :Including initial onboarding, on-going infrastructure monitoring, and 100 EPS
 - EPS Upgrade : Incremental 100 EPS for the remainder of term.
 - Temporary EPS Upgrade: Incremental 1000 EPS for a customer defined term.
 - Retention – 90 days default

Example:



Passport Advantage Parts

Current Parts	Description
D1SWCLL	IBM QRadar on Cloud 100 EPS Events Per Second per Monthly Subscription with Support
D1SWELL	IBM QRadar on Cloud 100 EPS Data Storage Events Per Second per Monthly Subscription with Support
D1GWKLL	IBM QRadar on Cloud 1K EPS Temporary Upgrade Events Per Second Monthly Subscription with Support
D1Q0WLL	IBM QRadar on Cloud Flows Add-On per 10K Flows per Minute Monthly Subscription with Support
D1Q0VLL	IBM QRadar on Cloud Vulnerability Management Add-On per 256 Assets Monthly Subscription with Support
D1PTLLL	IBM QRadar on Cloud Service Level Agreement
New Parts	Description
D1UCLLL	IBM QRadar on Cloud Log Archival 100 Events per second per Monthly Subscription with Support
D003TZX	IBM QRadar on Cloud Deployment Service Engagement
D003UZX	IBM QRadar on Cloud Optimization Service Engagement per Annum Subscription
D003SZX	IBM QRadar on Cloud Custom Parser Service Engagement

Overage Parts

D1SWDLL - IBM QRadar on Cloud 100 EPS Events Per Second Overage
D1SWFLL - IBM QRadar on Cloud 100 EPS Data Storage Events Per Second Overage



New QRadar on Cloud Parts – Details

D1UCLLL

Cloud Archival Parts - 100 EPS

- 100 – 100,000 EPS
- Extending cold storage for >3 months (or 1 full year) of active storage
- 2 requests per quarter; 30 days worth of data; 3 business day turnaround

D003TZX

IBM QRadar on Cloud Deployment Service Engagement

- Up to 40 hours of Product Professional Services
- Configure of events, activation of out-of-the-box rules, searches, graphs, and reports
- Custom tuning including the identification of removal of noise.

D003UZX

IBM QRadar on Cloud Optimization Service Engagement

- Provides ongoing reviews of Client's environments
- Addition of log sources, configuration of additional searches, reports
- Up to 8 days within the period of 1 year

D003SZX

IBM QRadar on Cloud Custom Parser Service Engagement

- Create, configure, and map a custom DSM
- Deploy and test the custom DSM



Sizing and Quoting

QRadar on Cloud Archival Parts

- Measured in Events Per Second (EPS)
- To quote multiple years, include additional quantity:
e.g. 1,000 EPS with 2 years of cold storage = 10 x 100 EPS x 2 (years)
- Quoted monthly; flexible billing options

Remotely Delivered Services – **Services parts not discountable**

Deployment Services

- 40 hours per part; no SOW required; expire within 90 days of purchase
- Includes an IBM Engagement Manager to schedule kick off calls and provide status updates

Cloud Optimization Services

- 8 days per year; minimum 2-day engagements

Customer Parser

- Provide the development of 1 custom parser (uDSM) for supporting Client's non-standard log source types to be sent to the Cloud Service
- Includes up to 25 message types for the log source



New Add On Parts - FAQ

- **What is the difference between the existing data capacity upgrade parts versus the log archive parts?**
- The Data Capacity Upgrade parts extend, active, searchable storage; the log archive parts provide cold storage. Cold storage must be re-mounted to the client's QRadar instance in order to be searchable.
- **For the deployment services, how many use cases and apps are included in an initial deployment?**
- The offer provides the implementation of up to ten use cases and up to two apps as offering time permits.
- **Do the Product Professional Services (PPS) parts for QRadar on Cloud require a Statement of Work?**
- No, the new Product Professional Services parts are available in Passport Advantage and do not require an SOW.
- **Are the PPS parts intended to provide ongoing managed services?**
- No, the parts are intended to provide initial, expert setup as well as ongoing tuning and optimization and not offence and alert escalation and management. The parts are complementary to add on managed services.
- **How to I make sure I am including the right amount of Services for a particular engagement?**
- If your client or partner is purchasing more than 7,500 EPS/50,000 flows and 4 Data Gateways, reach out directly to the Product Professional Team or Offering Management to find out how many multiples of each services may be required.



QRadar on Cloud – QRoC

The Data Gateway




The Data Gateway



- Customers must deploy data gateways to securely transmit security data to IBM QRoC
- Software is provided at no cost
- Customer has to provide its own Hardware or Virtual Machines
- Customer must have adequate bandwidth to send security data to IBM Security Intelligence on Cloud
 - **$\text{EPS_rate} * (\text{average event size} + 200) \text{ bytes} * 8 = \text{Mbps value}$**
 - Uplink is often either 10Mbps, 100Mbps or 1Gbps

What is a Data Gateway

- 15xx + qflow + vpn = Data Gateway
- Install on Bare Metal or VM
- Uses openvpn to connect to QRoC
- Buffers to disk if needed
- 10k eps or 200k fpm
- Does not currently support HA



Data Gateway

QRadar on Cloud – Data Gateway

The Data Gateway (DG) is a modified Event Collector transmitting data from the client's facilities to the Cloud via 4 key functions:



In the event of loss of connectivity, the DG will buffer to disk and transmit when connectivity is restored. The size of the buffer is client defined.

Deployment



Physical Appliance Specifications




CPU	2.6 GHz, 6 Core, 15 MB Cache
RAM	16 GB, 4 x 4 GB 1600 MHz RDIMM
HDD	2 TB:200 GB for software installation*

Virtual Appliance Specifications



CPU	4 cores for 1000 events per second (EPS) or less. 8 cores for 1000 -10,000 EPS.
RAM	16 GB, 4 x 4 GB 1600 MHz RDIMM
HDD	2 TB:300 GB for software installation*

QRadar on Cloud – Data Gateway



EPS and FPM limits for the QRadar on Cloud data gateway appliance

Events per second	Flows per minute
0	200,000
1,000	180,000
2,000	160,000
3,000	140,000
4,000	120,000
5,000	100,000
6,000	80,000
7,000	60,000
8,000	40,000
9,000	20,000
10,000	0

The Data Gateway - Adding

- The customer SaaS Admin accesses the cloud console and opens the Hosted QRadar icon
- There they will have access to:
 - Documentation
 - QRadar ISO
 - 7000 Appliance Activation Key
 - Gateway Token(s)
- Setup VM RHEL on VM or physical appliance 
- Install QRadar on top of RHEL using ISO from Hosted QRadar icon

Hosted QRadar

 Hosted QRadar  Manage Subscription

<https://console-demo01.console/adminconsole/jsp/SaaSGateway.jsp>

EULA ▼ Contact Support

Gateway Installation

[Installation guide](#)

Gateway Activation Key: 2Q3T7N-2V3A41-

QRadar Installation ISO: [qradar.iso](#)

Console Hostname: console-demo01

Gateway tokens are used during installation of a new data gateway, which collect log messages from your network for correlation and storage on the hosted solution. Once a data gateway is successfully connected to QRadar the gateway token will no longer be valid. Expired tokens are invalid.

For assistance acquiring a gateway token and deploying a new data gateway, contact SIaaSOp@ca.ibm.com.

Available Tokens


Host	Token	Expiration Date
No items to display		

Consumed/Expired Tokens

Host	Token	Status
	6566d9a2-0d70-	Retrieved Apr 5, 2016

The Data Gateway - Adding

- Select 7000 appliance key
- At the end of the normal setup the customer will get an additional prompt to add the gateway to the Console
- Connection configuration can be performed:
 - Automatically
 - Manually
- Gateway Tokens are valid for one use only
- For assistance deploying a new Data Gateway, contact: SIaaSOp@ca.ibm.com





Handling Data with QRadar on Cloud - QRoC






QRadar on Cloud– Types of Data Collected

- Events Generated from both on premise and cloud environments and synthesized with security data from cloud assets
- On-Premises flow data forwarded to the cloud
- On-Premises Vulnerability Scan Data forwarded to the cloud


QRoC – Collecting From On Premise

QRadar On Premise




QRoC – Collecting From On Premise and Cloud Example 1

QRadar On Premise



QRoC – Collecting From On Premise and Cloud Example 2


QRadar On Premise



QRoC – App Connectivity Patterns


On Premise

Data Sources
Applications
...




Cloud Services / Apps

Watson
Threat Intel
Resilient
...



QRoC – Disconnected Log Collector (Future)

QRadar On Premise



QRadar Vulnerability Manager for QRadars on Cloud

IBM QRadar on Cloud Vulnerability Manager proactively senses and discovers network device and application security vulnerabilities, adds context and supports the prioritization of remediation and mitigation activities.




Benefits

- **Fully integrated** with the IBM QRadar on Cloud Security Intelligence Platform
- **Sense and discover** network device and application security vulnerabilities
 - Coverage for over 70,000 known dangerous default settings, mis-configurations, software features and vendor flaws
- **Reduce critical exposures** and meet compliance needs
- **Use advanced IBM Sense Analytics™** to add context, identify key vulnerabilities and prioritize remediation activities
- Provide a consolidated vulnerability view across major vulnerability products and technologies

QRadar Vulnerability Manager for QRadar on Cloud – How it Works

- Activated with a licensing key and requires no new hardware or software appliances
- Deployed on your existing data gateway
- Asset discovery and vulnerability scanning
- Available in 256 asset increments





Flows for QRadar on Cloud

Flows for QRadar on Cloud provides flow analysis to help you sense, detect and respond to activities throughout your network.

Benefits

- **Fully integrated** with the QRadar on Cloud platform
- **Threat and anomaly detection** - Sense and detect new security threats without relying upon signatures.
- Gain **visibility** to malware, viruses and anomalies through behavior profiling for network traffic.
- **Advanced incident analysis and insight** - Perform near real-time comparisons of flow data (e.g. ports, addresses) with log events sent from security devices.



QRadar Flows for QRoC – How it Works

- The collector and the processor are deployed as software on your data gateway.
- Data is streamed to the hosted environment where it is available for correlation and display in the portal.
- The collector processes both internal and external flow data providing layer 7 and layer 3 network visibility.
- Supported flow sources include:
 - Qflow
 - Netflow
 - IPFIX
 - sFlow
 - J-Flow
 - Packeteer
- Available in 10K FPM increments

QRadar on Cloud– Supported Log Sources

- Consume Log and Service data from cloud based applications
 - E.g: AWS, Akamai, ZScaler
- Consume all other Qradar Supported log sources





QRadar on Cloud– DSM certificates

Contact sisaasop@ca.ibm.com if you require certificates for any of the following DSMs, or adapters to import certain data into QRadar.

- Amazon
- Generic Firewall
- Generic Auth Server
- IBM Endpoint Manager
- IBM Fiberlink
- Juniper Steel-Belted Radius
- Juniper Binary
- Open LDAP
- PostFix
- Salesforce Security Monitoring
- Sourcefire eStreamer
- Verdasys

QRadar on Cloud – Storing and Handling Data

- 90 Day Retention Base offering (it can be expanded per customer needs)
- Data at rest is encrypted
- Only Customer's provisioned users can access stored data – specified in the customer questionnaire that is filled out during the ordering process
- IBM DevOps and Operations team can access data only per customer's request
- IBM Operations team may assist setting up log sources only per customer's request





Ending QRadar on Cloud subscription

- If a customer decides to stop using IBM QRadar on Cloud, they must retrieve their data.
- To end the service, a customer must email sisasop@ca.ibm.com with information about when the service should be stopped
- IBM will send an email with the tokens that are required to stop the service, and instructions about how to retrieve the data.
- After a customer applies these tokens, they can no longer send events to IBM QRadar on Cloud.
- Customers are responsible for retrieving any data they want
- They have 30 days to retrieve any data that they want to keep.
- After 30 days all data will be expunged from IBM QRadar on Cloud



THANK YOU

FOLLOW US ON:



ibm.com/security



securityintelligence.com



xforce.ibmcloud.com



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



youtube/user/ibmsecuritysolutions

© Copyright IBM Corporation 2016. 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.



QRadar: Cloud Integrations



3rd Party Cloud Vendors



What Cloud?



Google Cloud Platform



Cloud Installs

Currently Supported



Planned



Google Cloud Platform



Cloud Ingestion

Currently Supported



Planned






Cloud Integrations



Securing cloud services and platforms



FAST VISIBILITY

Easily consume log and service data from cloud based applications

BUSINESS APPS

Comprehensive support for O365, Salesforce, Okta, etc.

CLOUD PLATFORMS

AWS, Azure, Soft layer, Open stack, zScaler, VMWare etc.

FIND THREATS IN THE CLOUD QUICKLY

Immediately discovers malicious activities in the cloud using existing analytics

Cloud Roadmap

Delivered	Q1 2018	Q2 2018	2H 2018
Integrations • AWS Security Content Pack • AWS CloudTrail • Azure Infrastructure Logs • Azure Event Hubs Hosting (BYOL) • AWS Software Install	Integrations • CloudWatch Logs and VPC Flowlogs • AWS GuardDuty • AWS Kinesis Streams • AWS Role Based Access for CloudTrail and Amazon Web Services Protocols • Cisco Umbrella • Azure Content Pack • Microsoft O365 Content Pack • AWS Security Content Pack Update Hosting (BYOL) • AWS Console AMI • AWS Managed Host AMI • Azure Console VMI • Azure Managed Host VMI	Integrations • QRadar AWS App • Amazon Inspector • AWS Config Rules • Azure Event Hubs Proxy Support • Amazon Macie Hosting (BYOL) • Google Cloud	Integrations • QRadar Azure App • Amazon WAF • Amazon SQS • Google Cloud Platform • Generic S3 Protocol • VPC Flowlogs in Network Activity Tab



AWS Deployment and Integration



Amazon integrations – S3 REST API protocol

- Initial support for Amazon event integration involved collecting CloudTrail events by downloading the log files from the S3 buckets where they were stored.



- “AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.”
- Originally this integration was implemented as a Service Type option in the Log File protocol source, along with the pre-existing SFTP, FTP and SCP options.



Amazon integrations – S3 REST API protocol

- In small-scale tests this approach worked but it ran into serious performance issues when dealing with large numbers of files. These performance limitations lead to the creation of a new protocol source specifically for interacting with the S3 storage service's REST APIs: "Amazon AWS S3 REST API"
- This specialized protocol scaled considerably better, and it has continued to be iterated upon.
- Support for Signature version 4 was introduced to stay up-to-date with the latest authentication protocol for the AWS APIs. AWS regions older than January 30 2014 still support signature version 2, but all newer regions require version 4.
- The currently available version of the protocol only allows connections to one region per log source, but an update to allow a single log source to retrieve logs from multiple regions will be released soon. Or we'll merge the S3 capabilities into our new "Amazon Web Services" protocol to benefit from it's ability to pull events from multiple regions.



Amazon integrations – S3 REST API protocol

- Although the initial use case for the protocol was focused around CloudTrail, the protocol can serve as a generic protocol source for retrieving any file from an S3 bucket
- We've already released the Cisco Cloud Web Security DSM, which uses the Amazon AWS S3 REST API with a W3C event handler to retrieve its files from S3.
- Cisco Umbrella will also make use of this protocol with a CSV event handler and will be released soon.
- We're also looking at adding a generic Line-by-line file handler for the protocol, so you can point it at any file in an S3 bucket and convert each line into an event.
- This protocol is available for all custom log source types, so once these latest updates are released, it will be possible to define a log source type (and log sources) for any set of events stored in S3 storage.



Amazon integrations – Amazon Web Services protocol

- We have another AWS protocol being released very soon: “Amazon Web Services”.
- The Amazon Web Services protocol is built on the AWS Java SDK, whereas the AWS S3 REST API protocol works through more general REST/HTTPS calls.
- It is currently in use by some customers as a beta – not long until general availability.
- In the long term this is meant to be a general protocol source for retrieving events from any Amazon service, though the initial release will be focused around connecting to the CloudWatch service.



Amazon integrations - Amazon Web Services protocol

- “Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.”
- Other Amazon services and applications can be configured to send their events/logs to CloudWatch, so it can serve as a single collection point, allowing QRadar to collect all desired events from CloudWatch after they have been collected there within the Amazon environment



Amazon integrations – Amazon Web Services protocol

- Within CloudWatch, logs/events are segmented into log groups, which can in turn be subdivided into log streams. A log source of the Amazon Web Services protocol can subscribe to a particular stream, or it can subscribe to all streams in a group.
- Each log source can only subscribe to one group, but like the most recent version of the REST API protocol, each can connect to multiple regions.
- It's in your best interest to organize the log groups and streams within CloudWatch is a way that plays nicely with how you want your log sources organized, so if that's possible, it's worth pursuing.

Add a log source

Log Source Name

Log Source Description

Log Source Type

Protocol Configuration

Log Source Identifier

Access Key ?

Secret Key ?

Regions ?

Asia Pacific Northeast-1

☐

Asia Pacific Northeast-2

☐

Asia Pacific South-1

☐

Asia Pacific Southeast-1

☐

Asia Pacific Southeast-2

☐

Canada Central-1

☐

China North-1

☐

Europe Central-1

☐

Europe West-1

☐

Europe West-2

☐

GovCloud

☐

South America East-1

☐

US East-1

☐

US East-2

☐

US West-1

☐

US West-2

☐

Other Region(s) ?

AWS Service ?

Log Group ?

Log Stream (Optional) ?

Filter Pattern (Optional) ?

Use As A Gateway Log Source ?

☒

Log Source Identifier Pattern ?

Use Proxy ?

☐

Automatically Acquire Server Certificate(s)

EPS Throttle ?

Enabled

☒

Credibility

Target Event Collector

Coalescing Events

☒

Store Event Payload

☒

Please select any groups you would like this log source to be a member of:

Save Cancel



Amazon integrations – Amazon Web Services protocol

- Because it's likely that CloudWatch will contain a heterogeneous mix of events from different log source types, the Amazon Web Services protocol was designed to serve as a gateway protocol. Like the various syslog protocol sources, a single log source using this protocol can feed events to multiple log sources.
- The protocol config parameters include a “Log Source Identifier Pattern” text area, which allows the user to list a set of <format string>=<regex> pairs. Each regex will be run against all events retrieved by the protocol; if one matches, the matching format string will be used to set the `sourceName` value on the resultant event payload object. The format string can use captured values from the regex

Amazon integrations – Amazon Web Services protocol

- As an example, the following two events are VPC Flow logs, obtained from CloudWatch:

```
{LogStreamName: eni-fa9996a8-all,Timestamp: 1508855283000,Message: 2
429269239926 eni-fa9996a8 77.72.82.14 172.31.25.226 40231 3471 6 1 40
1508855283 1508855342 REJECT OK,IngestionTime: 1508855428463,EventId:
33648597207639814566212319534420313607475850600323088387}
```

```
{LogStreamName: eni-7da14122-all,Timestamp: 1508855363000,Message: 2
429269239926 eni-7da14122 172.31.4.35 201.251.156.11 22 9224 6 16 3759
1508855363 1508855404 ACCEPT OK,IngestionTime: 1508855418101,EventId:
33648598991699430448662170844750398051943855577117229059}
```

- The “Message” field is the raw VPC flow log event; the other fields are added metadata from CloudWatch. The “LogStreamName” is the log stream (within a log group) where the events are stored.

Amazon integrations – Amazon Web Services protocol

- If you wanted to have a VPC Flow Logs log source for each log stream, you would configure the Log Source Identifier Pattern like so:

Log Source Identifier Pattern ?

```
VPC-$1=LogStreamName:\s(. *?),
```

- This would result in the preceding two events getting tagged with the following `sourceName` values, respectively:

VPC- eni-fa9996a8-all
VPC- eni-7da14122-all

- Because we are releasing a VPC Flow Logs DSM along with this new protocol, this means this would result in log sources autodetecting for the above two Log Source Identifier values (and for any other VPC instances with log streams in the target log group)




Amazon integrations – Amazon Web Services protocol

- Because the CloudTrail service can be configured to feed its audit events into CloudWatch, we can in fact use the new protocol to collect CloudTrail data, just as the REST API protocol can, though in this case the CloudWatch service acts as an intermediary.

Collect from the Cloud – AWS Infrastructure Logging


QRadar On Premise

AWS Ingestion Example




EC = Event Collector
EP = Event Processor
FC = Flow Collector
FP = Flow Processor

---> TLS



Collect from the Cloud – AWS Infra and Instance Logging Ex 1


QRadar On Premise



EC = Event Collector
EP = Event Processor
FC = Flow Collector
FP = Flow Processor

---> TLS


AWS Multi-Region Example



Collect from the Cloud – AWS Infra and Instance Logging Ex 2

QRadar On Premise


AWS Multi-Region Example



Collect from On Premise – Primary Infrastructure in AWS


QRadar On Premise

AWS Multi-Region Example



Collect from the Cloud – AWS Only Install

AWS Multi-Region Example





Frequently Asked Questions – AWS Integration

Q: Can QRadar retrieve the AWS CloudTrail logs from the root directory, such as /AWSLogs instead of /AWSLogs/<AccountNumber>/CloudTrail/<RegionName>/?

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

Q: Will AWS Role Based Access be supported?

A: We are targeting Q1 2018. This will allow QRadar ECs/EPs running in AWS to retrieve temporary credentials from their host EC2 instance, so you don't have to plug API keys into log source configs (and update them every time the keys need a refresh)

Q. What are VPC Flow logs?

A. VPC Flow logs capture information about the IP traffic going to and from network interfaces in your VPC (Virtual Private Cloud)

Q. What VPC Flow logs events are supported?

A. The traffic going in and out of your network interfaces in your Amazon VPC. Each event represents either an ACCEPT or REJECT – very firewall-like.



Frequently Asked Questions – QRadar deployment in AWS

Q: Is QRadar HA supported in AWS?

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

Q: In a hybrid deployment should I deploy an EC or an EP in AWS?

A: 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.

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

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

Q: When will the QRadar Amazon Marketplace Image (AMI) be available?

A: In Q1 we should have the QRadar console AMI on the marketplace. It will be bring your own license (BYOL) to start. The managed host AMI should follow shortly.



Azure Deployment and Integration





Microsoft Azure Event Hubs integration

- Late last year we released a DSM for Microsoft Azure
- It supports both syslog and a new protocol source, “Microsoft Azure Event Hubs”, which works very much like the new Amazon Web Services protocol, utilizing the Event Hubs Event Processor API to obtain Activity logs, Diagnostic logs, and Linux and other syslog messages that can be sent into Event Hubs.
- The Activity and Diagnostic logs are handled by our new Microsoft Azure DSM; any other events will be handled by existing DSMs as appropriate.




Microsoft Azure Event Hubs integration

- Naturally the Event Hubs protocol can serve as a gateway protocol as well, allowing retrieved syslog events to be routed to other log sources based on syslog header. For some reason it's missing a capability to customize how `sourceName` is set – I'll have to go yell at the integration dev team about that 😊
- The protocol requires both an Azure Storage account and an Event Hub Namespace and underlying Event Hub entity.
- Unfortunately there is currently no proxy support due to a limitation in Microsoft's SDK. They are working on it and expect to have it addressed in first quarter 2018, at which time we will update the protocol source to make use of it.

Collect from the Cloud – Azure Infra and Instance Logging

QRadar On Premise

Azure Ingestion Example






Microsoft Azure Event Hubs integration

- Microsoft also has an “Azure Log Integration” or “Azlog” service which is essentially an agent which can be installed on a Windows machine and will connect to the Azure cloud service and pull down event data. It can then forward the events via LEEF-formatted syslog to QRadar.

Collect from the Cloud – Azure Infrastructure and Instance Logging

QRadar On Premise


Azure Ingestion Example



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

QRadar On Premise

Azure Ingestion Example





Frequently Asked Questions – Azure Integration

Q. Is proxy supported?

A. No. AMQP (the message queuing protocol that Event Hubs components use) is a wire-level TCP protocol, it's not built on HTTP. So the Azure Event Hubs Event Processor library does not play nicely with web proxies. MS is working on it.

Q. Does the Azure Event Hub Protocol support Windows events?

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

Q. What is the retention period to store events?

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



Frequently Asked Questions – Azure Integration

Q. What kind of events can the protocol handle?

A. 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.
- Linux events collected from Event Hubs are received in a JSON wrapper but the raw events are extracted and treated like syslog so that auto discovery can figure out which Linux event type it falls under (DHCP server, iptables firewall or OS) and discover individual log sources appropriately.
- Generic syslog events are received in their raw form.



Frequently Asked Questions – QRadar deployment in Azure

Q: Can I install QRadar in Azure today?

A: Not at the moment.

Q: Why not?

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

Q: When is the marketplace presence coming?

A: We are targeting Q1 2017.

Q: Will QRadar HA be supported in Azure?

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



Installing in AWS



Installing QRadar in AWS - Today

- Choose your AMI: RHEL-7.3_HVM_GA-20161026-x86_64-1-Hourly2-GP2 from Community AMIs
- Choose your EC2 instance (M4.2XLarge or above based on [Virtual Appliance Sizing Guide](#))
- Choose 100GB for the root disk (GP2 is fine)
- 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
- Setup your security group to allow port 22 and 443 to a set of whitelisted IPs
- Choose your key pair or create one
- Review and Launch the Instance
- 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>:`
- As root run `aws_qradar_prep.sh --install`, then mount the ISO and run `/media/cdrom/setup`
- Use the internal IPs for the network configuration
- 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
```

- If you do this, there's no need to manually copy the iso or prep script to you instance, or to run the script. Already done for you!



Installing QRadar Community Edition - Today

- Choose the Centos 7 AMI from the AWS Marketplace
- Choose your EC2 instance (T2.Medium or above according to the [Community Edition Install Guide](#))
- 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)
- Setup your security group to allow port 22 and 443 to a set of whitelisted IPs
- Choose your key pair or create one
- Review and Launch the Instance
- As the centos user scp over the ISO:
 - Example: `scp -i <key.pem> QRadarCE7_3_0_20171013140512.GA.iso centos@<public ip>:`
- As root mount the ISO and run `/media/cdrom/setup`
- Use the internal IPs for the network configuration
- 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

- Choose the QRadar Console AMI or QRadar Managed Host AMI from the AWS Marketplace
- Choose your EC2 instance (M4.2XLarge or above based on [Virtual Appliance Sizing Guide](#))
- If it's a managed host enter the type of managed host in User Data
- Choose 100GB for the root disk (GP2 is fine)
- 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
- Setup your security group to allow port 22 and 443 to a set of whitelisted IPs
- Choose your key pair or create one
- Review and Launch
- 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" spifnolstp="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 --user
pip 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
```



THANK YOU

FOLLOW US ON:



ibm.com/security



securityintelligence.com



xforce.ibmcloud.com



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



youtube/user/ibmsecuritysolutions

© Copyright IBM Corporation 2016. 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.





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

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_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/>

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_client_vpn.html
- https://www.ibm.com/support/knowledgecenter/en/SS42VS_7.3.0/com.ibm.QRadar.doc/t_cloud_member_vpn.html