THE CLOUD FOR LAW ENFORCEMENT

WHAT YOU NEED TO KNOW NOW

MOTOROLA SOLUTIONS

1

WHY THE CLOUD AND WHY NOW?

Cloud computing has been around for many years in the enterprise and is gaining much more interest from government agencies for a variety of reasons, including cost savings, agility, and network simplification. However, you have unique requirements and data challenges. So how can cloud solutions best serve law enforcement and why should you transition to the cloud now?

While cost is an important factor, the most important benefit for your operations is the cloud's ability to drastically improve overall data management and utilization, ultimately helping you keep constituents safer.

THE DATA EXPLOSION

Data in all forms is exploding. According to Cisco, annual global IP traffic will reach 3.3 zettabytes (ZB) per year by 2021, or 278 exabytes (EB) per month. For reference, just in 2016, the annual run rate for global IP traffic was 1.2 ZB per year, or 96 EB per month (1 Zettabyte is the equivalent of 1 trillion Gigabytes). In this digital age, police are overwhelmed with structured data and, increasingly, unstructured data such as video, social media, sensor data and email. This tsunami of data is making it harder for you to be effective. For example, a typical search warrant nets approximately 3 terabytes of digital evidence data and without the proper tools it can take four weeks for investigators to go through just 1 terabyte.[1] Plus, they must process traditional structured data including arrest warrants, jail records, and criminal files.

The expanding Internet of Things (IoT) is also fueling a variety of enhanced public safety capabilities and generating even more data. So too is more and more video from body-worn cameras, in-car cameras, fixed surveillance and now even drones.

To make sense of it all, advanced analytics tools can help you extract the greatest value from this data. They help aggregate and correlate historical data and real-time information dispersed throughout your public safety operation. This is critical to enhancing public safety. Thus, there is now increased pressure to evolve your data management strategy and improve legacy systems and processes. Cloud solutions can and should play a critical role in this evolution.

BY THE NUMBERS

1 Zettabyte

=

1 Trillion Gigabytes

1 Zettabyte = 1000 Exabytes

1 Exabyte = 1000 Petabytes

1 Petabyte = 1000 Terabytes

1 Terabyte = 1000 Gigabytes

EVOLVING DATA MANAGEMENT SYSTEMS

Today, the majority of law enforcement agencies still rely on outdated, disparate systems that keep information in silos, reducing the ability to effectively derive and deliver intelligence where and when it's needed.

It's clear that police want better access to real-time, in-depth information to help focus resources on fighting crime smarter and more efficiently. Now, cities have started to embrace new systems and technology and they're seeing some important victories. For instance, the Chicago Police Department (CPD) is now widely deploying predictive and analytic tools after seeing positive initial results. Chicago Mayor Rahm Emanuel and police officials believe that using the latest in IT, including video surveillance and data-driven policing is reducing violent crime in the city.

Demographic changes in police forces will also affect the spread of data-driven technology. By 2020, millennials will make up 50% of the workforce.[2] This digitally-native generation is more tech-literate and skilled at multitasking. Agencies are already experiencing their influence as they rely on smartphone applications to assist with their responsibilities. Staying current with technology may help mitigate potential staffing challenges in the future and improve retention of qualified personnel.

Data volume and velocity will only continue to increase. Not only will the cloud greatly help to manage this explosion of data and provide powerful analytics tools, but because cloud services are an "additive" technology, not "rip and replace", departments can continue to get value from existing investments while amplifying their capabilities.

From managing massive amounts of data to recruitment of millennials who expect the latest technology tools, the cloud is pivotal to the needs of the public safety community as it continues to evolve.

> **Demographic changes in police forces will affect the spread of data-driven technology. By 2020, millennials will make up 50% of the workforce[2]**

2

WHAT IS THE CLOUD?

The cloud leverages a different approach to buying, managing, and deploying IT infrastructure and software solutions. With private, on-premise solutions, applications are deployed by the customer, on customer-owned equipment, and subsequently managed by the customer. Customers are responsible for everything including service availability, data durability, data security, geo-redundancy, scaling, and more.

Conversely, hosted cloud solutions are offered as-a-service, deployed in a private, community or public cloud at a cloud service provider's data center. In this instance, the infrastructure is managed by the cloud service provider and can be operated solely for an organization or small group of organizations (private or community cloud), or made available to the general public or larger, typically industry-oriented group (public cloud).

Lastly, hybrid solutions can be a combination of hosted cloud and private, on-premise solutions.

TECHNOLOGY DEPLOYMENT MODELS

CLOUD

Cloud Service Provider
Hosted and Fully-Managed

ON-PREMISE

Customer Fully-Managed
Data Center

HYBRID

Combination Of Cloud
and On-Premise

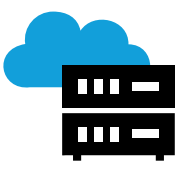

CLOUD DEPLOYMENT MODELS

PRIVATE	The cloud infrastructure is operated solely for an organization by a cloud service provider, typically as isolated infrastructure within their data center.
COMMUNITY	The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns. (e.g.,mission, security requirements, policy, and compliance considerations).
PUBLIC	The cloud infrastructure is a made available to the general public or a large industry group.

CLOUD OFFERINGS

When dealing with hosted cloud solutions (whether in a private, community or public cloud environment), there are different models for interacting with and consuming cloud resources. Generally speaking, these are classified as: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS).

IaaS provides you with the least amount of resources from a cloud service provider. Customers purchase and consume infrastructure including computing, networking, and storage resources. This is the most flexible model but also requires the most work on the part of the consumer, as customers are purchasing basic capabilities, and deploying their own solution onto the IaaS resources. These customers are fully responsible for the performance of such deployments.

Moving up the stack, PaaS allows cloud consumers to also purchase services from a cloud service provider rather than just purchasing basic resources as in IaaS. Services vary by provider, but may include databases, container management tools, messaging, API gateways, and load balancing. With PaaS, the basic components of your solution are being provided. That said, since the services and interfaces vary among cloud service providers, applications constructed using PaaS resources are much less portable between cloud service providers. Even with the simplified application development provided by PaaS, a significant amount of work is still required by the cloud consumer to create complete solutions.

SaaS presents cloud consumers with complete software solutions in the cloud. Customers of SaaS solutions purchase or lease application resources to directly solve a problem. A familiar consumer example is Netflix, a SaaS solution provider which provides their customers with a complete video streaming solution built around applications for streaming, billing, account management, etc. To put this in perspective, Netflix is a consumer of PaaS and IaaS from Amazon Web Services (AWS), using these resources and services to build applications and deploy a SaaS solution, which their customers lease from them, paying a monthly fee to access.

For most state and local public safety agencies SaaS solutions will offer the most comprehensive answers to their policing needs.

CLOUD CONSUMPTION MODELS

= You Manage = Others Manage

ON-PREMISES	INFRASTRUCTURE AS A SERVICE	PLATFORM AS A SERVICE	SOFTWARE AS A SERVICE
Applications	Applications	Applications	Applications
Data	Data	Data	Data
Runtime	Runtime	Runtime	Runtime
Middleware	Middleware	Middleware	Middleware
O/S	O/S	O/S	O/S
Virtualization	Virtualization	Virtualization	Virtualization
Servers	Servers	Servers	Servers
Storage	Storage	Storage	Storage
Networking	Networking	Networking	Networking

3

SECURING DATA IN THE CLOUD

When considering the security of different technologies, it's very important to understand that the greatest threats to an organization's systems have less to do with where they are located, but rather how they are accessed, how data is protected at rest and in transit, and how the applications are secured.

Most attacks don't target the infrastructure, but rather the application layer via injection attacks, broken authentication, insufficient access control and security misconfiguration, all coupled with an inadequate level of logging and monitoring. These are all areas where cloud service providers in conjunction with SaaS solution providers delivering applications, excel. While on-premise solutions could theoretically do the same, the IT security budget of cloud service providers dwarfs even the largest department IT security budget. Still, security is about risk mitigation and cost-benefit analysis and there are both positive and negative security considerations for different deployment models.

The greatest threats to an organization's systems have less to do with where they are located, but rather how they are accessed, how data is protected at rest and in transit, and how the applications are secured.

ON-PREMISE

There are three primary areas where on-premise security excels:

1. When there is a need to be air-gapped, or otherwise not connected to the public Internet
2. When there is a data privacy regulation or other compliance need for data to reside in an area not supported by the cloud
3. When absolute control over security is desired

However, control does not necessarily translate to security. Properly securing the applications, data, services and infrastructure you own comes at a considerable cost, both in terms of capital expenditure as well as staffing IT security resources to configure and maintain the security tools and the security posture. This may be time and money that could be better spent focusing on the core mission of your organization.

CLOUD

With a SaaS deployment in the cloud most of the security is handled for you by a cloud service provider (e.g. AWS) and SaaS solution provider (e.g. Netflix). This frees precious IT security resources to focus on who should have access to data. Tooling provided by these providers enables simple point and click security administration to control access to data and services and also provides simple interfaces for auditing and notifying administrators when security configurations are modified. Such administrative tools make security misconfiguration practically a thing of the past.

Cloud is also particularly valuable for its cyber-resilience and ability to fend off targeted and sustained Distributed Denial of Service (DDoS) attacks such as those launched by Botnets. While cloud detractors would suggest that the concentration of information in the cloud makes it a tempting target for bad actors, only the sheer size and scale of the cloud has the resources to adequately handle large scale attacks. Cloud providers go to great lengths to protect against attack. It would be nearly impossible to match that level of defense on-premise.

Threat Intelligence is another area where cloud excels. Cloud service providers have extremely sophisticated monitoring tools and may very well be the first to notice an attempted attack or breach immediately alerting the solution provider or even yourself. In a world where bad actors collaborate and coordinate, it is extremely beneficial to have a cloud service provider working on your behalf. Additionally, because these providers see so many attacks against so many of their tenants, they are the first to apply security best practices based on the type of attack patterns they notice. In doing so, an organization benefits from the lessons learned across a wide sample of other organizations.

Cloud service providers and SaaS solution providers also remove much of expensive and time-consuming burden of pursuing security compliance certifications, such as DoD certifications, FIPS certifications, ISO certifications, and other country-specific requirements. Such certifications and attestations not only give your IT department confidence that the solution is secure, but they also save your IT department time and resources, freeing budget to be spent elsewhere.

HYBRID

For those that want to leverage many of the security benefits of the cloud, but still want to maintain a tighter degree of control over their data, a hybrid approach might make sense. Hybrid-cloud deployments give the customer varying degrees of control, such as allowing the customer to control data encryption keys that still exist in the cloud, all the way to maintaining the encryption keys on-premise and encrypting data on-premise before shipping to the cloud. In the latter example, all data in the cloud is encrypted in transit and at rest using encryption keys that never leave the customer's premise.

Hybrid security can take on other forms as well. For instance, customers can continue to provision and credential their employees on premise, then use those same identities and credentials to access a SaaS solution. This enables the reuse of existing Identity & Access Management solutions and allows a single place to on-board and off-board users regardless of whether they access a solution on-premise or in the cloud. This is just one example of hybrid cloud combinations. Many types are possible, with the customer maintaining full control over what exists on premise and a large degree of control over the parts that exist in the cloud.

4

SAAS SOLUTION CONSIDERATIONS

When choosing a cloud-based SaaS solution, there a number of issues to weigh. Security, both physical and cyber. Data privacy. Service availability. Scalability. Redundancy. Where to start? Below, we discuss some of the most important factors when deciding on a new SaaS solution.

SECURITY AND PRIVACY

While SaaS solutions generally incorporate strong security capabilities, it's important to verify them and ensure you're comfortable with their competencies. To that end, when making an assessment, both the security of the underlying cloud service provider as well as the security of the SaaS solution provider needs to be taken into account.

For the cloud service provider, the physical security of their data center is a critical attack vector and needs to be considered. It's important to understand how server rooms are secured, who has access, and how visitors are vetted. Is physical access to the server rooms logged, monitored and audited? In the case of government operations, are authorized personnel vetted citizens of the country? Are background checks performed on administrators and others with access to resources? It's also critical to understand where cloud servers are physically located and if you need the data to physically reside in a specific country or physical area. And, what geographic redundancy requirements exist?

Security in the cloud, just like on-premise, is all about how data is handled, including its confidentiality both in transit and at rest, its integrity, and its availability. You must understand and be comfortable with who has access to the data, when they have access to it, and from where they have that access. Every access to data, whether it be to read data, write data, create data, or delete data, should be logged. Standards are also mandatory here. One of the golden rules in security is "don't invent your own." A SaaS solution provider should be capable of explaining the usage of vetted, industry-based, open standards for security.

> **One of the golden rules in security is "don't invent your own."**

These additional questions should also be considered:

Encryption

Is data encrypted in transit and at rest using open standards encryption?

Patch Management

Are known vulnerabilities actively patched in a timely manner?

Threat Intelligence

Does the solution provider alert you if your data has been compromised or attempted unauthorized access?

Application Security

Is the code continuously scanned for known vulnerabilities?

Compliance

What criteria is the solution provider compliant to?

Bring Your Own Credential

Does the solution provider let you leverage your existing Identity & Access Management infrastructure?

Multi-factor Authentication (MFA)

Does the solution provider offer secure and usable MFA technologies that are resistant to phishing attacks?

Policy Enforcement

Does the solution provider enable you to enforce your IT security policies?

DDoS Attacks

Does the combination of the cloud service provider/SaaS solution provider actively defend against Distributed Denial of Service (DDoS) attacks, ensuring that access to your data and services remains available even in the face of a sustained and targeted attack?

Assessment

Has the solution provider's security been assessed by an independent third party?

Knowing that your data is protected from bad actors is a good start, but it also must be protected from misuse by the solution provider as well. **As the owner of the data, you should be entitled to ownership over the data and how it's used.** It is important to ensure that there are policies in place by your cloud service provider and SaaS solution provider that prevent inappropriate access of data. Providers should also only hire vetted employees and administrators who have passed extensive background checks and are recommended with multiple, credible references. In addition, if the decision is made to leave the SaaS solution provider, you should get to keep your data. Data sovereignty should be a mandatory requirement.

AVAILABILITY AND SCALABILITY

AVAILABILITY

Availability describes a system or service's ability to continue operation in the presence of hardware and software failures. It's a function of the mean time between failure (MTBF) and mean time to repair (MTTR) reliability parameters. It's generally represented as a percentage or a fraction, for example .9995, which indicates the percentage of time the service is expected to be available. High availability is desirable for services due to the importance they play in your operations, whether they are operating in a public or private cloud environment or on dedicated hardware on-premise.

Another aspect of availability is performance guarantees, which can be considered a type of gradual or soft failure. A service that can't meet its performance guarantees cannot be considered highly available. For this reason, consider how you can manage availability both in the presence of failures and in the presence of increased workloads.

> **High availability is desirable for services due to the importance they play in your operations.**

Maintaining high availability for a service requires addressing the following issues:

1. Ensuring that the service will continue to operate in the presence of limited hardware failures
2. Ensuring that the service will continue to operate in the presence of limited system software failures
3. Ensuring that the system continues to perform as expected in the presence of increasing workload

The first two of these require some level of redundancy to address and the third requires a degree of system scalability.

One system resource required for both on-premise and cloud solutions is the presence of networking infrastructure. For SaaS solutions, this network is a combination of cloud service provider network infrastructure, telecom provider connectivity, and the Internet. For on-premise solutions, this network consists of customer-owned network infrastructure. The best practice for achieving highly available networking with a SaaS solution is to use redundant hardware and telecom providers on both sides of the network connection to the cloud service provider. To achieve high-availability connectivity, connection hardware must be redundant, even when connecting from the same location. Cloud service providers maintain this type of redundancy in their systems. The same approach is true for on-premise connectivity, although all redundancy must be maintained by the owner of the on-premise solution.

The Internet itself is another point of failure, but the Internet was designed redundantly with failures in mind. Since the Internet is simply the interconnection of multiple public and private networks such as Comcast, Verizon, and universities, and since these connections are redundant and geographically diverse, there is little concern that the Internet will be unavailable. The constituent networks that make up the Internet are also deployed with redundancy built-in, but even so, they may fail, which is why it's recommended to connect redundantly through multiple telecom providers. Increasingly, on-premise solutions are requiring Internet connectivity to interface with mobile devices. So, in many cases being on-premise only partially protects you from Internet connectivity failures. Still, making use of the Internet using appropriate levels of redundancy is almost certainly more reliable than a private network maintained by a single entity.

Making use of the Internet using appropriate levels of redundancy is almost certainly more reliable than a private network maintained by a single entity.

In order to ensure availability, the following principles should be followed:

1. All cloud deployed solutions need to use cloud-specific exclusion or availability rules such that single failures in the underlying hardware will not render the solution unavailable.
2. All solutions need to be deployed using management and orchestration software that deploys them across Virtual Machines (VM) such that single VM failures will not render the solution unavailable.
3. All VM failures should be detected and new VMs brought on-line automatically.
4. All solutions need to be deployed using management and orchestration software such that interruptions are detected and the affected services are automatically restarted.
5. All solutions need to be monitored for performance degradation and failure so additional instances of the solution may be brought online to increase service performance or replace failed services. This includes solutions that are deployed redundantly for high availability.

SCALABILITY

Scalability can be classified as either horizontal or vertical. Horizontal scalability of a service allows for additional copies of its components to be added or removed to account for different workloads. Vertical scaling is the technique of adding processing power, bandwidth, and storage to an existing service's underlying hardware without adding additional copies of service components.

For horizontal scalability to be achievable, services must be explicitly designed for scalability. This scalability may be applied at two different levels of the system hierarchy – scaling within an availability zone (AZ), and scaling across multiple AZs. AZs are physically separated processing resources and can be either "geographically" distinct deployment locations within a single cloud service provider's system, or separate cloud service providers. Scaling across availability zone has the added advantage of increasing availability. Scaling within an AZ also increases availability in as much as the service will degrade with individual component failures rather than failing all together.

One additional step can be taken, but is not recommended. This is deploying a solution to multiple cloud providers such as AWS and Google. The additional cost of developing, deploying, and maintaining two different instances of each service provides little benefit over and above the benefits already accrued via individual cloud provider redundancy and AZs.

In order to ensure scalability, the following principles should be followed:

1 Solutions need to be designed to support single AZ scalability such that additional instances of the components that comprise the solution can be added or removed without affecting functionality such as stateless components, load-balancing, data storage and partitioning.

2 Solutions need to be designed to support multiple AZ scalability such that additional instances of the solution, each in a separate AZ, can be added or removed such as inter-AZ load balancing and data synchronization.

3 Cloud infrastructure on which solutions are running needs to utilize monitoring software so overload situations can be detected and additional compute and storage resources may be deployed to handle the increased load.

4 Solutions themselves also need to be monitored so that additional resources may be brought on-line in the presence of increased workloads.

MONITORING

Monitoring is another critical aspect of operating solutions in the cloud. In deployments that may consist of hundreds or thousands of resources including servers, databases, and cloud-based applications, effective monitoring allows you to have insight into the status, health, and general trends that exist across these resources.

Monitoring presents three primary challenges: determining which resources and metrics should be monitored, deciding how this data will be collected into a monitoring tool and how error alerts are handled.

Cloud service providers typically provide basic monitoring tools that can be used to collect and visualize data generated by their services. APIs published by these vendors allow SaaS solution providers to export their own internal data into the cloud monitoring platform as well. However, these public cloud monitoring services are generally basic and must be augmented with some combination of commercial or open source tooling.

Once logs and metrics have been identified and tooling is in place to collect the data, it is then important to understand how much of this data you will collect and at what interval the collection will occur. Collecting a snapshot of CPU utilization every 30 seconds, for example, results in almost no impact to a running system but will result in a less accurate view of near-real-time performance than if it were sampled every 100 milliseconds. Log data, in particular, can result in extremely large data sets, particularly when debug messages are stored or in the event of a large number of error conditions. Cloud solutions are naturally better equipped to handle this massive amount of data but impact to running systems, cost of storage, and granularity of metrics being collected must always be weighed to make sure the benefits outweigh these costs.

The final, and perhaps most challenging, aspect of any successful monitoring implementation is the process of alerting operations personnel or customers to error conditions. The ideal alerting configuration filters out false positives and duplicate messages while notifying recipients of all valid error conditions so they can respond effectively in a timely manner. This should be a foundational focus area of any organization launching products or services in the cloud.

Logs, events, and metrics can be aggregated to provide a picture of the current state of software and physical resources. However, the ultimate measure of successful delivery and availability is the end user's actual ability to use an application. SaaS solution providers typically offer availability guarantees to customers in the form of Service Level Agreements (SLAs). SLAs generally specify the expected availability of solutions usually expressed in the form of a percentage, as explained above, with guaranteed credits or even refunds in the event the provider does not meet the SLA in a given time period. Effective monitoring helps support SLA reporting by providing an audit trail of successful operations.

ADVANTAGES BY MODEL

	ON-PREMISE	CLOUD	HYBRID
Cost	• On-premises deployments are typically structured as a one-time capital expenditure, eliminating the need for recurring monthly costs. • Owned hardware (servers) can be virtualized and shared for other internal needs at the owner's discretion.	• SaaS solution costs are classified as OPEX, reducing the need for approval of large CAPEX during a given year. • Cloud storage costs continue to decrease. Extra storage can be added with no hardware to purchase. • Support / maintenance costs are included in monthly costs and not as an additional fee, making budgeting easier.	Hybrid solutions, depending on their configuration, often include many of the benefits from both on-premise and cloud solutions, plus the following: • Hybrid allows a less expensive way to add capabilities to an existing, on-premise system.
Security	• Security can be highly customized to an individual organization's processes, requirements and regulatory requirements. • Knowledge of system and data reside solely in-house. May be preferable for agencies with security needs where they do not want data "shuffled"' over the Internet or want to restrict Internet access to their systems, databases or applications entirely.	• Physical security, infrastructure security, and application security is handled for you. • A higher degree of cyber-resilience helps fend off targeted and sustained Distributed Denial of Service (DDoS) attacks, keeping services available when you need them most. • More out-of-the-box certifications mean you don't have to worry compliance such as for DoD, FIPS, ISO and other country-specific ones. • Cloud service providers actively and successfully monitor against threats on your behalf and alert you to concerns.	• Various security controls can be put in place on-premise to accommodate specific agency needs that the cloud does not provide for.
Deployment and Scalability		• Cloud solutions can typically be deployed in a matter of days (weeks/months for more customization) because hardware and software does not have to be installed onsite. • Cloud solutions are highly scalable. Organizations can simply request more seats or storage and attain it rapidly. • Software updates and bug fixes can be done more frequently with the ability to more easily roll back any breaking changes. • Cloud solutions typically require less IT involvement and less in-house technical skill for deployment, updates and changes.	• On-premises infrastructure can support average workloads with the ability to leverage the cloud for failover circumstances in which the workload exceeds the power of the on-premise resources. • Critical data from an on-premise solution can be replicated to a cloud environment in a different location to the primary systems to ensure business continuity in event of catastrophic failure.
User Accessibility	• On-premises systems can run without Internet access. This is useful to keep mission-critical apps running - particularly in areas where Internet connectivity is not reliable.	• Cloud solutions require Internet connectivity. With wired and wireless broadband become cheaper and available virtually anywhere, it provides easy access for remote workers via multiple devices.	

5

KEEPING COMMUNITIES SAFER

Police departments face many challenges, from an explosion of data, to mobility and IoT, to the recruitment, retention and training of officers. Technology should help alleviate challenges and make officers more effective without adding new burdens. But outdated, disparate systems that keep information in silos, can reduce the ability to effectively derive and deliver intelligence the way you need it.

To be sure cloud-based SaaS solutions are right for you and your department, it's important to understand how this technology fits and improves on your existing systems and how to responsibly select the right vendor. By following a framework that helps you properly weigh security, availability, scalability, and system monitoring you can be confident in selecting the right cloud solution.

Indeed, forward-looking agencies are realizing that cloud-based SaaS solutions, properly vetted and intelligently selected, can drastically improve data management and data utilization. With these solutions in place, data transforms from a burden to being the engine powering the data-led policing strategies that keep our communities and officers safer.

REFERENCES

1. "Digital age a turning point for policing, says commissioner Leppard," ComputerWeekly, May 14, 2015, http://www.computerweekly.com/news/4500246279/Digital-age-a-turning-point-for-policing-says-commissioner-Leppard

2. "PwC: Millennials at Work" https://www.pwc.com/gx/en/managing-tomorrows-people/future-of-work/assets/reshaping-the-workplace.pdf

MOTOROLA SOLUTIONS

Motorola Solutions, Inc. 500 West Monroe Street, Chicago, Il 60661 U.S.A. **motorolasolutions.com**