Launchmetrics Security Policy and Practices

Document Status:	**Version 2.1**
Approved by:	**Pau Montero Parés - CIO - 03/2018**

Record of Policy Review and Updates 3
Notices and disclaimer 3

Launchmetrics Security Policy and Practices 4
Application Security 4
Backup and Disaster Recovery 5
Hosting 6
Logging 6
Server Access 7
Support 7
Security 8
Data Architecture 8
Office Security 9
Incident Management 11
Monitoring 12
Account & Password Management 12
Appendix: AWS Certifications and Security Compliance Resources 14
Global Certifications 14
European Certifications: 14
US Certifications: 14
Security Compliance Resources 14

Record of Policy Review and Updates

Date of Review	***Reason for Review***	***Lead Reviewer***	***Version***
2013/09	Creation of First Security Policy	Anish Singh	Version 1.0
2013/10	Update Policy	Mario Soave	Version 1.1
2014/03	Added References section	Anish Singh	Version 1.2
2016/04	Datacenter changes	A Merad	Version 1.4
2017/06	Added data center certifications	Ozgur Ozdemircili	Version 1.5
2018/01	Office network security	Jordi Guerrero	Version 1.6
2018/02	Security Enchantments	Ozgur Ozdemircili	Version 1.7
2018/03	Detailed informations for several sections	Ozgur Ozdemircili	Version 2.0
2018/03	Proofread revision	Marine Leclinche	Version 2.1

Notices and disclaimer

This document is provided for informational purposes only. It represents Launchmetrics' current and broad security policy and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and together with any use of Launchmetrics' products or services. This document does not create any warranties, representations, contractual commitments, conditions or assurances from Launchmetrics, or its affiliates. The responsibilities and liabilities of Launchmetrics to its Customers are controlled by Launchmetrics' Service Provider Agreements, and this document is not part of, nor does it modify, any Agreement between Launchmetrics and its Customers.

Launchmetrics Security Policy and Practices

At Launchmetrics we treat data and application security very seriously. Providing excellent services like strong security, high SLA and prompt resolution of client issues are key priorities driving our day to day activities as well as long term strategy. Separate policies have been adopted for application security, data and information management procedure, Disaster recovery, Incident management, Change Control and Asset management. Each of these documents can be provided to you upon request.

Below is a high level summary of our procedures in several critical areas concerning data and information security.

1. Application Security

1.1. We have regularly scheduled security audits, both internal and external.

1.2. Frequent external security audits and vulnerability scans are performed by TrustWave. The results of these audits are available to customers on request.

1.3. We use Secure Sockets Layer (SSL/TLS).

1.4. We abide by strong cryptographic standards, including advanced password hashing techniques such as sha256 mixed with crc32b.

1.5. Strong incident management, change control and asset management policies. The process is detailed in Section 10.

1.6. Access to our applications may be restricted only to whitelisted IP addresses.

1.7. Password Authentication for all users: only authorized users have access to the application.

1.8. Support for different roles and permissions for each role. Permissions can be set at the role level or at individual users level. Only roles or user authorized to access a protected resource can do so.

1.9. All User activity is logged: in the event of unauthorized activity, we can review the log to investigate the events and provide the log to the client if requested.

1.10. We do use one time password authentication for critical systems. AWS, Gmail, Github, Lastpass applications are all secured with the second layer of OTP system where the user is required to input username and password as well as the code shown on the authenticator application.

1.11. All databases used by the backend systems are placed in individual networks seperated for production and staging environments. Production servers are only accessible by production SRE team. All actions on these databases (import / export / changes) are tracked and logged.

1.12. No real client data is allowed in the staging environment. All client data is fake and re-created anonymously for testing purposes. Backups and restores are done automatically without human intervention. Every import into staging or development environments is done after data sanitization and anonymization.

1.13. All user information along with username / email addresses and contact numbers are deleted from the user database.

1.14. All information that need to be shared with the client are shared either over an Sftp channel or via an encrypted bucket.

1.15. Part of the applications or the products working in relation to provide service to our clients are communicated via a) VPN b) HTTPS c) Encrypted ports.

1.16. All client data is separated from other in logical level. Only the client with the correct password is allowed to decrypt and access to its own application and data.

1.17. All application changes to be deployed to production environment are audited and tested by our QA team. First tests are completed in Staging environment, necessary changes are communicated to the development team where they are corrected. Then the production deploy QA team and the development team together with the SRE team work to make sure the correct deployment has been done, there is no impact on overall system and database changes are applied correctly.

1.18. Data Exchanges between different Launchmetrics applications are encrypted via HTTPS.

1.19. All our application and servers are patched regularly.

2. Backup and Disaster Recovery

2.1. All our databases are backed up every day automatically. e are firstly doing backups without affecting the main database performance by creating slave databases and we keep a real time copy of the databases in case of disaster. The database backups are realized in two distinct ways to ensure quality and recovery:

2.1.1. SQL dumps: Every 12 hours. A logical copy is made and uploaded to an S3 distributed bucket. This backup is encrypted and saved. The data retention period is 6 months.

2.1.2. Incremental binary backups. Allow data recovery in any point of time. The data retention period is 31 days.

2.2. All our application servers are secured and distributed behind load balancers. We are able to detect the traffic and do maintenance in the servers without affecting the client service.

2.3. Our Database servers are created with Multi-AZ option which gives the ability to have master-slave configurations running in different datacenters in the same location. In case of any disaster in any of the datacenters, we are able to switch to the slave database without losing any data.

2.4. Our Application and Database servers are spread over three different availability zones in two different regions. Keeping service alive in case of failure of one availability zone is automatic. In case of losing a region, service continuity is also secured.

2.5. All S3 backups are encrypted.

2.6. Restoration process is done automatically in a daily basis to ensure the quality of the backups.

3. Hosting

3.1. We do hold certifications*(*1)* in appendix assuring the security and integrity of our servers and data.

3.2. Depending on the location that client chooses the client application and databases are hosted in AWS North Virginia, or Ireland locations.

4. Logging

4.1. Types of logging:

4.1.1. Server Access (network address, user, date and time)

4.1.2. Infrastructure modifications via Cloudtrail (network address, user, server/service, date and time)

4.1.3. Account activity (user, modifications, date and time)

4.1.4. Application and services activity

4.1.5. Inter-application activity

4.2. We do continuously collect streams of real time information from all our servers as well as services (i.e. VPN, Access to AWS Services, Lambda executions). These information are collected under one Cloud Trail Dashboard where we can see all the data streaming as well as do a point-in-time trace to investigate what happened during a period of time.

4.3. All trail logs are checked and audited by production staff and security alerts are put in place for failed logins, data usage and unauthorized access. All logs are date / time stamped.

4.4. All user activations, deactivations, file usage and configuration changes are logged and audited with Cloudtrail log streams.

4.5. All streams are archived and saved for 90 days.

4.6. These logs are only accessible with production SRE team and no client is allowed full or partial access to these logs. Only authorized and audited logs are provided to clients that match their own activity. No client is allowed access to other client’s logs.

4.7. In case of security audit and for investigation purposes, Launchmetrics can provide access to log activity related to the investigation.

5. Server Access

5.1. Our AWS administration employees require an account, a password and a One Time Password that is automatically created every 30 seconds using multi factor authentication. After 3 wrong attempts, users are blocked and can only be unlocked via our SRE team.

5.2. Server access is managed via Public Key Authentication with SSH protocol. Access via password authentication is disabled by default. Key pairs are asked to be rotated every 90 days making sure that even if there is an intrusion the old key pairs are unusable.

5.3. Staging and Production servers are running in different subnets to allow restricted access to authorised users.

5.4. Individual server access rules or role based are applied for all personal. All users are required to connect to our VPN server in order to be able to connect to any servers. Once connected a role is assigned that allows restricted access to the designated services. This way even though the user has the private key for a different user, he will not able to login to any server that he is not permitted to access.

5.5. Only our production SRE team is allowed to access production application and data servers. Our Software Development Directors are only allowed production application access. They are not authorized to access servers.

6. Support

Launchmetrics do hold an enterprise level support contract with our provider Amazon Web Services which gives access to the following:

6.1. 24/7 available Technical account manager

6.2. AWS Trusted Advisor Dashboard: Prevent attacks and provide information on the overall security of out servers

6.3. AWS Personal Dashboard: Personalized view of AWS services and alerts

6.4. Well Architected Review: Continuous Overall review for having secure, efficient, reliable infrastructure.

6.5. 4x7 access to Sr. Cloud Support Engineer via email, chat, and phone with the following SLA:

6.5.1. General guidance: < 24 hours

6.5.2. System impaired: < 12 hours

6.5.3. Production system impaired: < 4 hours

6.5.4. Production system down: < 1 hour

6.5.5. Business-critical system down: < 15 minutes

7. Security

Servers are hosted at two Data Centers: Amazon AWS East US (Virginia), AWS EU Ireland.

7.1. AWS provides state-of-the-art security measures in the following layers.

- **Perimeter layer:** AWS data center physical security begins at the Perimeter Layer. This Layer includes a number of security features depending on the location, such as security guards, fencing, security feeds, intrusion detection technology, and other security measures.

- **Infrastructure layer**: The Infrastructure Layer is the data center building and the equipment and systems that keep it running. Components like back-up power equipment, the HVAC system, and fire suppression equipment are all part of the Infrastructure Layer. These devices and systems help protect servers and ultimately your data.

- **Data layer:** The Data Layer is the most critical point of protection because it is the only area that holds customer data. Protection begins by restricting access and maintaining a separation of privilege for each layer. In addition, we deploy threat detection devices, video surveillance and system protocols, further safeguarding this layer.

- **Environmental layer:** The Environmental Layer is dedicated to environmental considerations from site selection and construction to operations and sustainability. AWS carefully chooses our data center locations to mitigate environmental risk, such as flooding, extreme weather, and seismic activity.

7.2. You can find detailed information about AWS servers security here: AWS Security Compliance Resources

7.3. Launchmetrics software and client data is hosted in AWS owned and maintained by Launchmetrics. Images are stored on a Content Delivery Network, an integral service by AWS, to provide faster download times from users all over the world.

7.4. All our servers are protected against security issues using the following methods:

- **Port Security:** All services exposed to clients are accessible only by load balancers. Load balancers are configured to expose only necessary ports creating a shield for the other server ports that may be open.
- **DDOS Security:** All load balancers are setup to block any repeated petitions per

second.

- **Load Balancer Security:** automatically distributes incoming application traffic across multiple targets, such as Amazon Elastic Compute Cloud (Amazon EC2) instances, containers, and IP addresses, and multiple Availability Zones, which minimizes the risk of overloading a single resource. Elastic Load Balancing, like CloudFront, only supports valid TCP requests, so DDoS attacks such as UDP and SYN floods are not able to reach EC2 instances. It also offers a single point of management and can serve as a line of defense between the internet and your backend, private EC2 instances.
- **Network Security:** All servers are maintained in different physical networks (VPCs), and in some cases in different accounts disabling any intentional or unintentional communication request between servers and products. Between these networks exist firewalls that only let the configured ports to be open to communications from one specific ip to another.
- **User Security:** All staff are assigned an IAM user. All IAM users belong to one or if necessary more groups which has specific roles associated to it. Without specific roles associated staff are unable to enter / use specific services / servers.
- **DMZ Zone Security:** All our production servers and only front-end servers are placed in a seperate network where they are only accessible behind a load balancer. This way we do create a DMZ network where we secure the rest of our infrastructure in case of any security breach.
- **DNS Security:** We use Route53 service to handle DNS requests.Amazon Route 53 is a highly available and scalable DNS service designed to route end users to infrastructure running inside or outside of AWS. Route 53 makes it possible to manage traffic globally through a variety of routing types, and provides out-of-the-box shuffle sharding and Anycast routing capabilities to protect domain names from DNS-based DDoS attacks.
- **HTTP Static Content Security:** Amazon Cloud Front distributes traffic across multiple edge locations and filters requests to ensure that only valid HTTP(S) requests will be forwarded to backend hosts. CloudFront also supports geoblocking, which you can use to prevent requests from particular geographic locations from being served.
- **Security Monitoring:** All security related incidents are audited by our SRE team on daily basis. Cloud Trail logs are also saved for future references. Using the Guard Duty and Trusted Advisor tools we are able to see the overall security problems and any suspicious connection attempts from and to our servers are services.

7.5. All passwords, documents related to access and passwords are kept in our cloud provider. Launchmetrics employees are required to comply Clean Desk Policy leaving no client data visible or on their tables.

7.6 Rather than storing the encryption key in a local file, we use AWS KMS where possible. For example: we start serverless queries which will connect to AWS Key Management Service and ask it to generate a new key. Lambda keeps the key on-disk in an encrypted form until the the job finished. The key stored on-disk cannot be used to decrypt the data; rather, on each startup, Lambda connects to AWS KMS and has the service decrypt the locally-stored key(s).

The decrypted key is stored in-memory as long as the process is running, and that in-memory decrypted key is used to encrypt the local data.

7.7 The IAM keys are paired with an individual key which is audited via our systems such as Trusted Advisor Board. We are alerted with the necessary renewal of an existing key as well as their last login time. All keys are rotated every 90 days.

8. Data Architecture

8.1. Each client has only the permission to access to their own databases and tables. Each client has its own set of users, groups and permission levels. Access is not shared between clients.

8.2. Our data servers are running highly available, multi-zone replication configurations with automatic failover mechanisms which gives us the ability to switch to the new master in case of a datacenter-wide disaster.

8.3. Our application servers are also distributed in multiple zones in the same region. In case of a disaster in one of the datacenters we will continue providing service using the rest of the application servers without any interruption to our clients

8.4. All out backups are saved in encrypted S3 buckets

9. Office Security

9.1. All connections to AWS servers are maintained by VPN connections and secured by various methods. All users in all of our offices are cable-connected to a main switch where they are grouped depending on their tasks. Our main switch assigns a special IP address to our employees which is granted only for the environments where they are authorized to connect.

9.2. For outside VPN users, we use the same method but the n assigned rights are based on usernames.

9.3. All our offices networks are centrally controlled in a cloud environment where we have a general view of the traffic, problems and attacks.

9.4. We do have physical firewalls in the offices as well as on AWS premises to avoid any attacks.

9.5. Documents are saved in Teamdrive, a are shared space where teams can easily store, search, and access their files anywhere, from any device and securely. Files in Team

Drive belong to the team instead of an individual. Even if a member leave, the files stay exactly where they are so your team can continue to share information and get work done. Documents are shared only among authorized employees. Documents on the office network are not public and can only be accessed by authorized employees or consultants.

9.6. Access to the building is not granted unless the visitor is pre-authorized or a current employee.

9.7. Security center network: the Security Center provides a centralized view for security filtering events. This includes both IDS/IPS and Advanced Malware Protection (AMP) events. It provides information and insights to a network administrator through a variety of different components, each focusing on different analytics and uses.
Components list:

9.7.1. Summary View: the summary view of the Security Center provides a variety of visual components to understand the security events in the network.

9.7.2. Retrospective Malware Detections: this component provides alerts about downloaded files that have changed to a malicious disposition.

9.7.3. Events over time: the Events over time component shows the number of events matching configured filters, over a specified interval of time, ranging from one month to two hours.

9.7.4. Most affected clients: this component provides a breakdown of the subset of clients that have generated the most events for the selected filters.

9.7.5. Top sources of threats: this component provides both a map and a table summary of the most common IP addresses associated with threats matching the configured filters. The map provides a visual view into the trajectory of these threats, from the network location to the geo-located source of the IP address associated with the threat.

9.7.6. Most prevalent threats: this component provides a list of the most frequent threats matching the selected filters. These can be the most common IDS/IPS signatures that have been detected, the most frequently scanned or blocked file through the AMP engine, or a combination of both.

9.7.7. Most affected operating systems: this component summarises the events matching the selected filters by client operating system. The events are aggregated based on the operating system of the client devices in the security events and are displayed in the table by the number of events associated with that operating system.

9.7.8. Events View: the Events view provides the same data as the summary view in a text-based log. It is still possible to filter this data in the same ways as the summary view.

9.8. Threat protection is comprised of the Sourcefire® SNORT® intrusion detection engine and AMP anti-malware technology.

9.9. Advanced Malware Protection (AMP): Advanced Malware Prevention inspects HTTP file downloads through an MX Security Appliance and blocks or allows file downloads based on threat intelligence retrieved from the AMP cloud. When traffic is filtered, the URL or ID and the action taken are logged in the Security Center. Advanced Malware Protection (AMP) is an industry-leading anti-malware technology from SourceFIRE, integrated into. MX Security Appliances.

9.10. Intrusion detection and prevention: intrusion detection feeds all packets flowing between the LAN and Internet interfaces and in-between VLANs through the SNORT® intrusion detection engine and logs the generated alerts to the Security Report. You can also export these alerts via Syslog. Intrusion prevention blocks traffic that is identified as malicious, rather than just generating alerts for it.

9.11. Meraki Firewalls: all offices with meraki firewalls have two very important extra network controls

9.11.1. Traffic shaping: includes an integrated Layer 7 packet inspection engine, enabling QoS policies, load balancing, and prioritization based on traffic types and applications.

9.11.2. Content filtering: allows to block certain categories of websites based on the organization policies. Can also block or whitelist (allow) individual websites for additional customization.

9.12. Apple Macbook hard drive encryption: all new laptops purchased in the company have to be encrypted for security with FileVault full-disk encryption. FileVault uses XTS-AES-128 encryption with a 256-bit key to help prevent unauthorized access to the information on your startup disk. During 2018 we hope to have the entire fleet of encrypted laptops. Encryption occurs in the background and only while the Mac is awake and plugged in to AC power. FileVault requires that you login every time your Mac starts up, and no account is permitted to log in automatically.

9.13. Two step verification in Google accounts: Two Factor Authentication, also known as 2FA, two step verification or TFA (as an acronym), is an extra layer of security that is known as "multi factor authentication" that requires not only a password and username but also something that only, and only, that user has on them, i.e. a piece of information only they should know or have immediately to hand - such as a physical token. The token that can be received it by text message to a mobile phone or using Google Authenticator: a software token that implements two-step verification services using the Time-based One-time Password Algorithm and HMAC-based One-time Password Algorithm, for authenticating users of mobile applications by Google.

10. Incident Management

10.1. The Fireline is the process to handle all technical incidences in the company. This includes: customer requests, security issues, bugs, platform issues and questions.

10.2. Detailed description of the Fireline is in the document Technical/Engineering Incident Management: Fireline Process

10.3. Our SRE team works in weekly changing shifts where they are expected to ACK an incident within 15 minutes. The detailed information of the On-call duty organization is described in SRE On Call Process

10.4. In case of a security problem (vulnerability) a special group is created to work on the impact of the problem. A remote conference call is started with the software developers, hosting partners and the technical teams to do an in-depth analysis of the situation. These results are added to a security case document to provide quick resolution. The resolution is applied to all necessary systems to prevent further leakage.

10.5. After a security incident if any client data is breached or if any sensitive information is made public client, a notification of the situation along with the details of the incident and the known information leak is sent to the Customer without undue delay..

10.6. After a security incident and when appropriate, Launchmetrics will to communicate the detailed report of incident to the competent authority.

10.7. If the Customer have any security question, it can reach out to its Service Provider's contact or to security@launchmetrics.com

10.8. For each security incident detected, an internal investigation will start to understand the cause and origin of the breach and find the person responsible, if any. If the person responsible is found then he or she will be reported both internally and to the appropriate authorities without undue delay. If the person responsible is an employee of the company, our executive team will take the appropriate disciplinary action. All security breaches are investigated to find the root cause and a responsible. In case the responsible is public then he/she is reported to the necessary authorities, in case the responsible is a member of the company a disciplinary action is taken.

11. Monitoring

11.1. All applications entry points are monitored on 18 different locations 24/7. The application server metrics are monitored using a seperate monitoring system which enables detection of process fails, connection and service problems. Network connection speeds, server cpu, memory and disk IO usages, service uptimes are monitored and connected to the main alerting system.

11.2. All applications are monitored using a monitoring user that is seperate to clients' user and password.

11.3. All Cloud Trail, Trusted advisor, Guard Duty, Monitoring alerts are continuously monitored by SRE team as well as AWS Enterprise support technical staff.

12. Account & Password Management

12.1. TFA is integrated on all main sensitive services. All administration services that are being used require input of a user and password as well as a OTP that is created every 30 seconds. Unless all of these are input correctly, the user is not permitted the access to the required resource.

12.2. Launchmetrics does embrace SSO technology to authenticate users. SSO technology, along with TFA gives us the advantage of permitting or revoking any access to any resource in our systems.

12.3. Customer account users are isolated from each other by application. Password security is manage by encryption and encryption rules depend on the application. Password enforcement rules are also dependant on the application. Description provided below.

12.4. Password rules for all applications

12.4.1. Samples, Events, Contacts.

12.4.1.1. Password must contain at least 8 characters.

12.4.1.2. Password must NOT be created based on the username.

12.4.1.3. Password must contain at least one character in lowercase.

12.4.1.4. Password must contain at least one character in uppercase.

12.4.1.5. Password must contain at least one number.

12.4.1.6. Password must contain at least one special character (! % & @ # $ ˆ * ? _ ~)."

12.5. Users, groups, roles and permissions are defined in AWS IAM.

12.6. Only SRE team members have read-write access to client data. These group managed by SRE Director reporting to security Officer and CIO of the company. Software Development Directors are provided read data to production systems for debugging purposes.

13. Data Elimination or Return

Upon the expiration or earlier termination of client contract or a special request by client, and at the choice of client:

13.1. All client data is securely return to client or its designee OR;

13.2. All client data is eliminated in the following way by the SRE team unless Launchmetrics is required by Applicable Data Protection Laws to store the Personal Data:

13.2.1. All users that were operated by the client are eliminated from the application in coordination with the development directors.

13.2.2. All databases created for the client is removed from RDS instances.

13.2.3. All client related data is deleted from company registers.

13.2.4. All backup data related to the client is eliminated from related S3 buckets.

Appendix: AWS Certifications and Security Compliance Resources

Global Certifications

ISO 9001

ISO 27001
ISO 27017
ISO 27018
PCI DSS Level 1
SOC 1 : Audit Control Reports
SOC 2 : Security, Availability And Confidentiality Report
SOC 3 : General Controls Report

European Certifications:

Austria TUV
Germany C5
Germany IT-Grundshutz
UK Cyber Essentials Plus
UK G-Cloud

US Certifications:

Criminal Justice Information Services
DoD Data Processing
FedRamps Government Data Standards
Ferpa Educational Privacy Act
FFIEC Federal Financial Institutions Examination Council
FIPS Government Security Standards
FISMA Federal Information Security Management
GXP Quality Guidelines and Regulations
HIPAA Protected Health Information
ITAR International Arms Regulations
MPAA Protected Media Content
NIST National Institute of Standards and Technology
Sec Rule 17a-4(f) Financial Data Standards
VPAT / Section 508 Accessibility Standards

Security Compliance Resources

https://aws.amazon.com/compliance/resources/