PROCURE SECURE
Continuous monitoring for public sector cloud services

Dr. Giles Hogben

European Network and Information Security Agency


APPROVED
100% PURE
100%
BIO
CERTIFIED
NATURAL PRODUCT
100% ORAGANIC


FARM FRESH
FARM FRESH

Continuous monitoring
The proof of the pudding is in the eating

enisa
European Network and Information Security Agency

Cloud Computing

Information Assurance Framework

November 09

enisa
European Network and Information Security Agency

What is continuous monitoring?

- Real-time service level data/feeds, including service level dashboards.
- Regular service level reports.
- Incident reports and alerts raised by the cloud provider.
- APIs

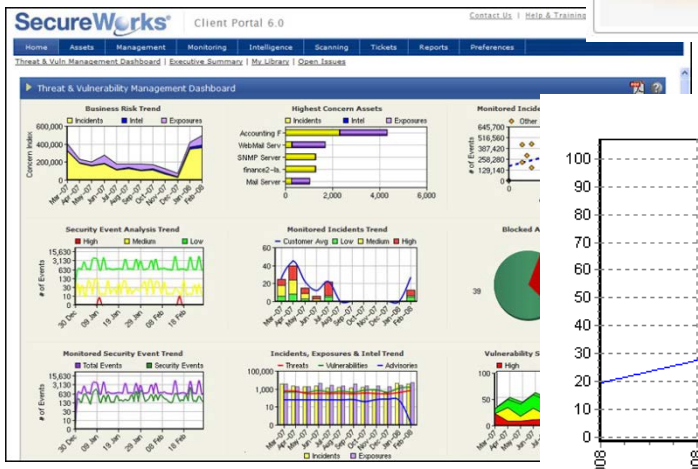

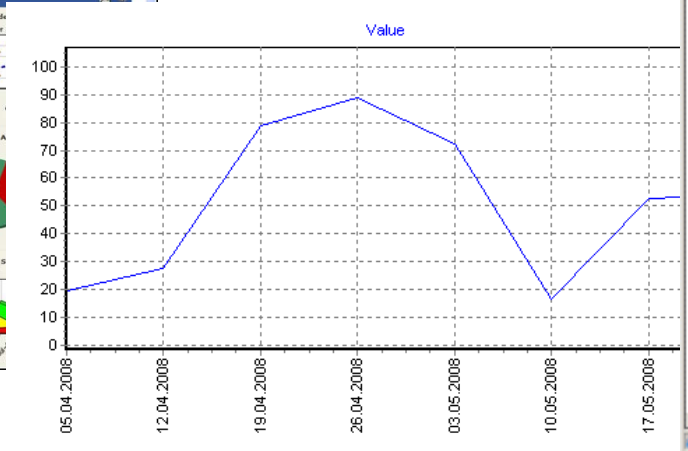

http://is.gd/fwDwgf

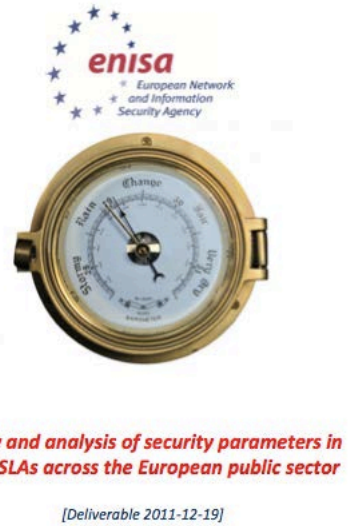

ENISA SURVEY ON CONTINUOUS MONITORING IN THE PUBLIC SECTOR

Survey and analysis of security parameters in cloud SLAs across the European public sector

- 117 fully completed responses from IT officers across the European public sector
- 15 different EU countries
- 77% of respondents said they have high or very high security requirements (41% and 36%)
- 70 respondents agreed to be part of the focus group.

Penetration tests

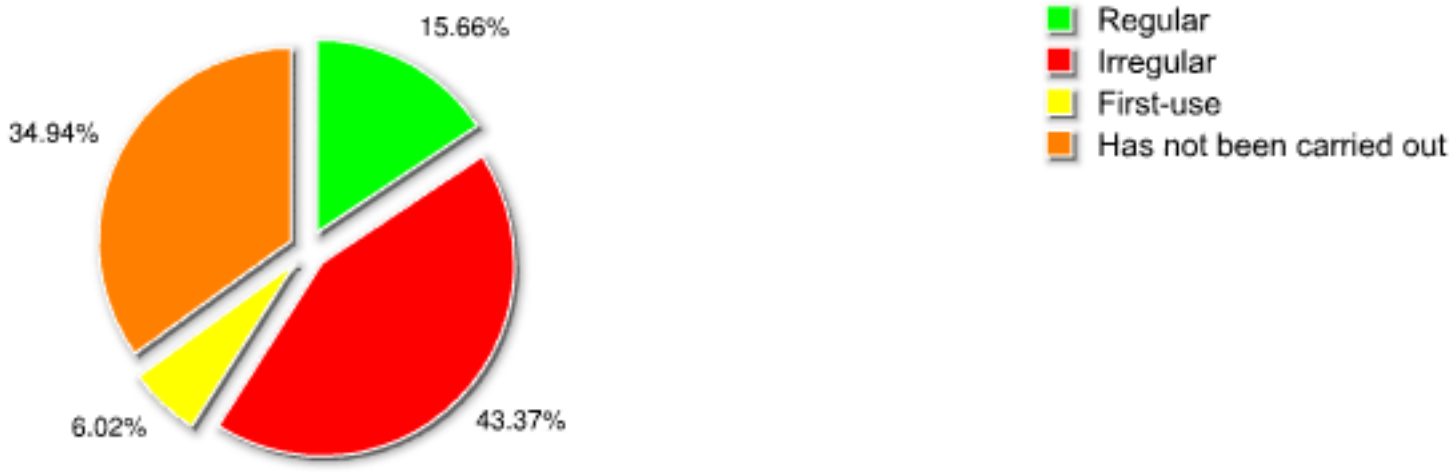

Backup/failover tests

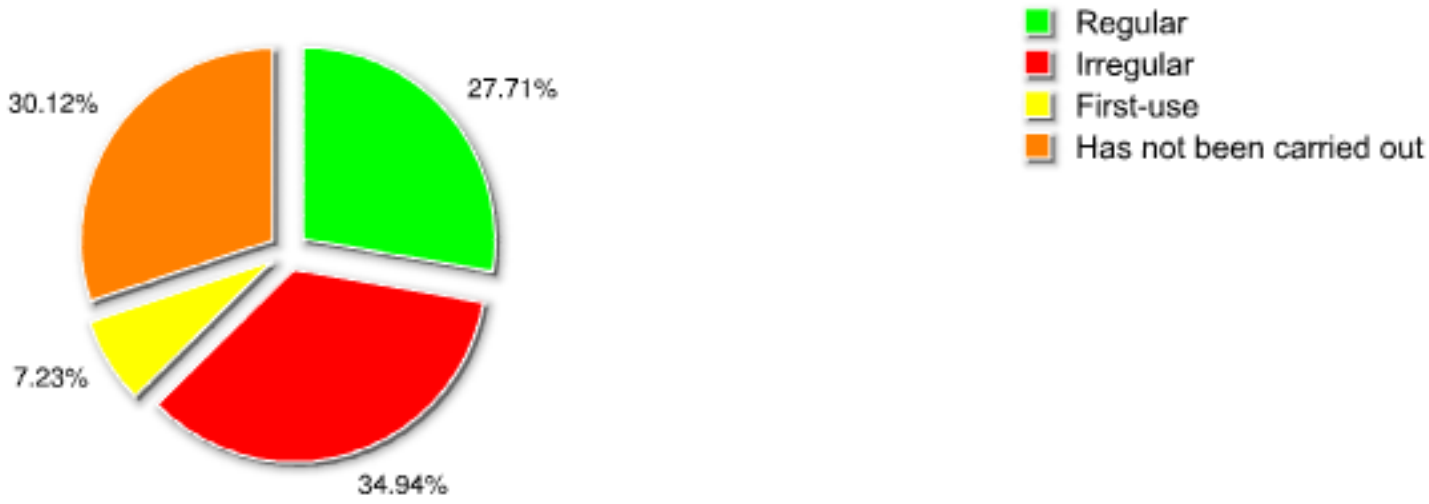

Data portability tests

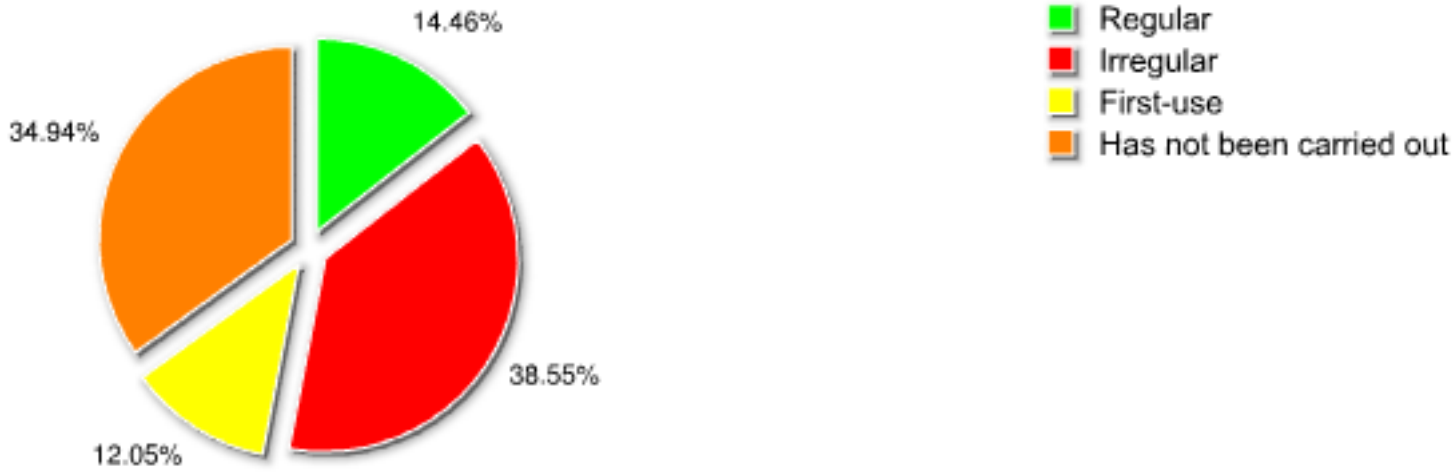

http://is.gd/syMAjD

MAIN REPORT

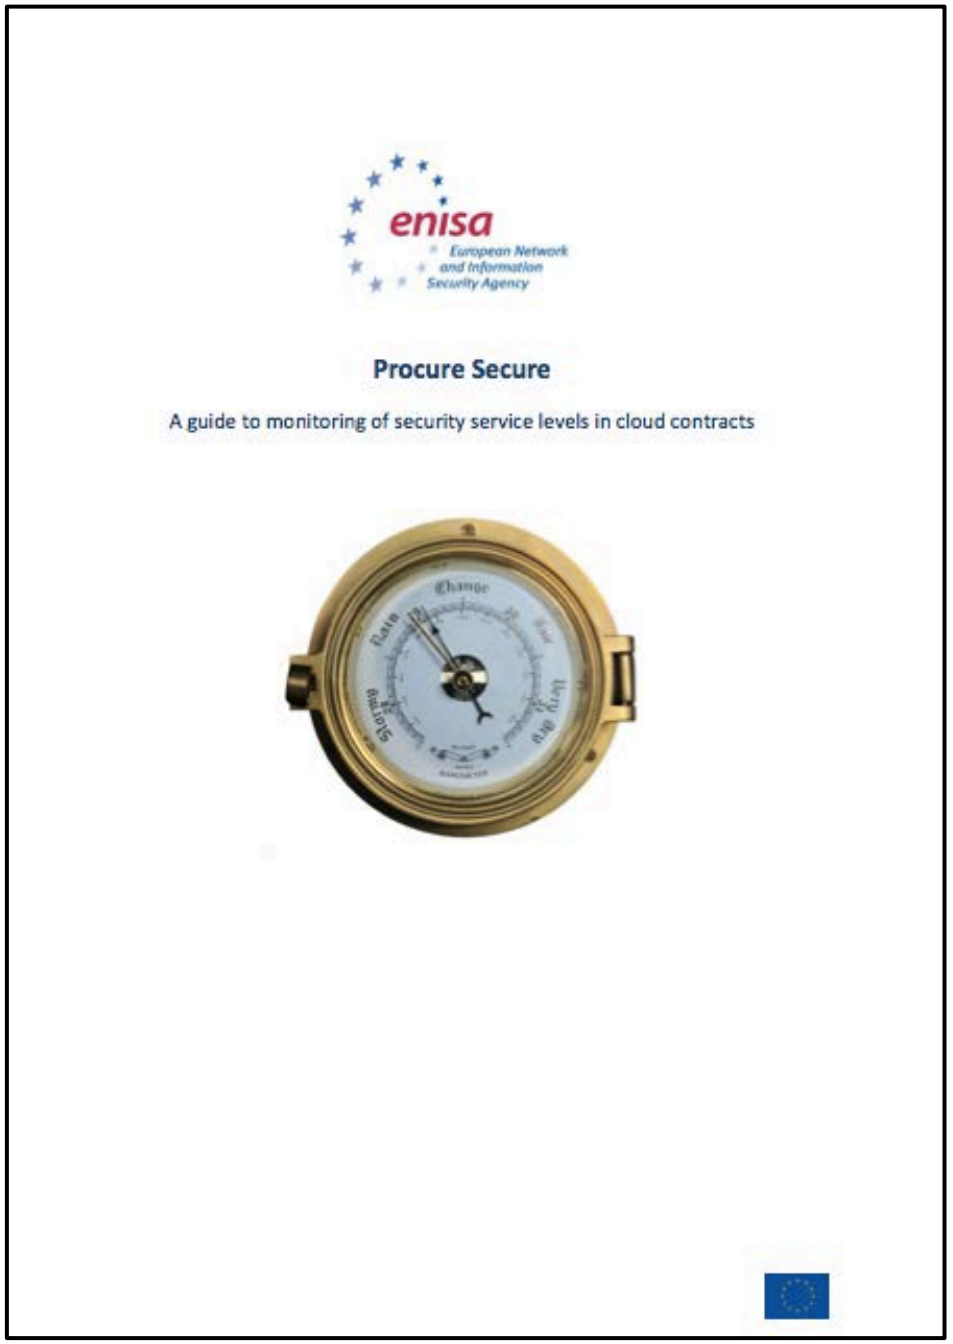

Who contributed?

- Paolo Balboni, ICT Legal Consulting, Tilburg University, European Privacy Association
- Art Barnes, Dell Secureworks
- Matt Broda, Oneforo Corporation
- James Bryce Clark, OASIS
- Daniele Catteddu, Cloud Security Alliance
- George Chetcuti, Government of Malta
- Nick Coleman, IBM
- Dr. Peter Dickman, Google
- Dr. Niels Fallenbeck, Fraunhofer AISEC
- Julia Herman, European Aviation Safety Agency
- Brian Honan, BH Consulting
- Jens Jensen, Science and Technology Facilities Council, UK, Funded by EU Contrail Project
- Ben Katsumi, IPA, Japan
- Kieran McCorry, Hewlett Packard
- Mark Penny, UK Department of Health Informatics Directorate
- David Pollington, Microsoft
- James Reynolds, Left Coast Logic
- Dobromir Todorov, Amazon Web Services
- Dr. Nicolas Vanderavero, Federal Public Service Finance, Belgium
- Beau Woods, Dell Secureworks

The parameters

1. Service availability
2. Incident response
3. Service elasticity and load tolerance
4. Data life-cycle management
5. Vulnerability management
6. Change management
7. Data isolation
8. Log management and forensics

Parameter breakdown

- What to measure
- Should I care about it
- How to measure it
- Independent testing
- When to raise the flag/thresholds
- Customer responsibilities

Drill-down

- Service availability
- Incident response
- Service elasticity and load tolerance

Availability

- What to measure?
  - Scope: How many users?
  - Scope: What service functions?
  - Define failure: When is a user “available”
  - Commitment period: Can I have all my unavailability in one go? Does it matter more at weekends/nights
  - Scheduled unavailability
- Do I care
  - E.g. Scheduled unavailability at weekends
  - Large transactions and MTBF

Availability

- How to measure it
  - User reports
  - Logs: Examination of logs by the provider, to detect errors.
  - Sample requests/service health-check.
- Independent testing
  - Polling, user feedback (make sure you don't trigger DDoS protection, or CAPTCHAs).

Availability

- When to raise the flag/thresholds
  - How realtime is your service – e.g. financial services would set much lower thresholds for availability incidents.
- Customer responsibilities
  - Understand dependencies
  - For systems under your control (e.g. IaaS servers)
    - Design for failure where you can
    - Test and monitor.

Examples

Example 4: As part of a disaster recovery plan, a government provisions a number of virtual servers from an IaaS provider to maintain public services. It maintains them on continuous 'warm' stand-by so that they can be used at short notice when required. By agreement with the CSP, the servers are polled every 10 minutes by an automated system. If any of the servers do not respond, the customer administrator is notified by SMS.

Example 6: For a customer using an IaaS provider to run web servers, if the percentage of all http requests [*request definition*] to servers for all the provider's customers, as registered in logs shared by the provider [*scope of the service*], with a specified error status, is more than 5% [*failure of a service request*] for at least 5 minutes [*sample size*], then the server is considered to be unavailable (in those 5 minutes). This assumes that the provider is willing to publish the availability figures recorded in their logs.

Incident Management

What to measure

- What is a severe incident
- How many severe incidents have occurred and how quickly did the provider respond?
- What % of sev x incidents are resolved within time y.
- Does the provider keep you up to date?
- How quickly do they detect (where there's an independent measure).

How to measure

- Incident classification scheme

Incident management

- Independent testing
  - Independent logs of response times
  - Independent detection can tell you about detection times (or failure to detect)
- Customer responsibilities
  - Make sure you're not causing an incident
  - Agree on classification scheme
  - Provide any customer-side resources required to resolve an incident

Example

Example 12: An SLA for an online tax return service states the following objectives:
The mean time to respond to an incident will be under 2 hours for a given month. A severity classification scheme detailing levels from 1 to 5 is defined in the agreement. A monthly service level report includes data on:

- Average time to respond to customer-reported incidents.
- The number of severity 3 and above incidents resolved (system functionality restored) within 1 day.
- The number of exercises performed to test response procedures.

Incident alerts are provided to the customer via encrypted email, based on triggers defined in the SLA:

- When time to respond to customer-reported incidents rises above 4 hours.
- When an incident occurs with severity level of 4 or above.

Where serious intrusions are detected, the CSP will report:

- Assets affected
- Data breaches detected (number of records or data classification, to the extent this is known)

In the case of a severity 4 or above incident, the provider will report on a 2 hourly basis measures taken to resolve the incident and an estimated time to recover.

Elasticity and load tolerance

- What to measure
  - Ratio of failed resource provisioning requests to total number of resource provisioning requests
- Should I care
  - Load volatility
  - DDoS risk
- How to measure
  - Burst testing
  - Real-time monitoring or log inspection of resource provisioning

Elasticity and Load Tolerance

- Independent testing
  - Depends on overall demand -> independent testing is often meaningless
  - Test reserved capacity limits and provisioning speed

Examples

Example 16: An online tax return service has 100 virtual servers in constant active use. In order to cope with periods of peak demand, it has a further 400 virtual servers reserved for immediate availability and a contracted ability to grow their provision by 200 virtual servers per day, daily for up to one week. Thus they could grow to 1800 virtual servers in a week. It would be sensible for them to test this, by both moving some of the 400 reserved virtual servers into active use and by growing the reserved pool (within the 200 VM/day growth limit). If any of those requests failed this should raise an incident report. Such a test might cost the customer a small amount of money, but is a sensible validation of their provider's ability to deliver on the contracted service.

Example 17: The service provider of the online tax return service provides regular reports detailing the ratio of provisioning failures to the overall capacity in operation, as recorded in their logs.

Howard Schmidt, Whitehouse Cyber-Security Coordinator

- **Continuous Monitoring of Federal Information Systems**

“Transforms the otherwise static security control assessment … into a dynamic risk mitigation program that provides .. near real-time security status and remediation”

Neelie Kroes, Davos, Switzerland, 26th January 2012

"Today I am inviting public authorities and industry, Cloud buyers and suppliers, to come together in a European Cloud Partnership.

In the first phase, the Partnership will come up with common requirements for Cloud procurement. For this it will look at standards; it will look at security;"

- Commission has proposed to allocate 10 million Euro in funding for common procurement requirements in 2013

Procure Secure – Continuous Monitoring

- Anyone procuring IT systems
- Focus on cloud, public sector but widely applicable.
- If you are busy- use the checklist format

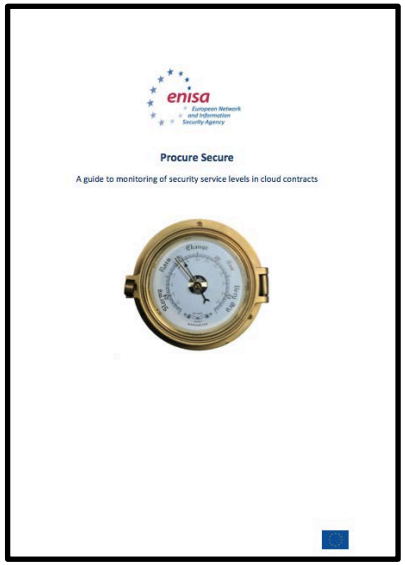

- Survey: http://is.gd/fwDwgf
- Guide: http://is.gd/syMAjD