Cradle-7
From concept to creation...

3SL

Cradle Capabilities for Information Assurance

RA003/07 June 2019

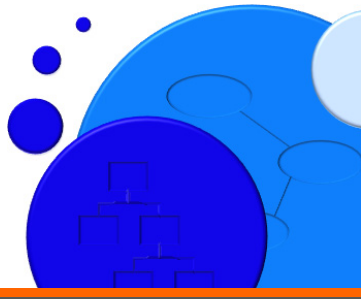

http://www.threesl.com
salesdetails@threesl.com
support@threesl.com

Structured Software Systems Ltd (3SL)
Suite 2, 22a Duke Street
Barrow-in-Furness
Cumbria LA14 1HH, UK
Tel: +44 (0) 1229 838867
Fax: +44 (0) 1229 870096
Regd: 2153654 VAT: GB 473 2757 28

Contents

Introduction

Any organisation achieves its goals by doing work, using data that must be protected to:

- Safeguard its intellectual property
- Hold all personal information securely
- Protect military or government secrets
- Safeguard national security

An organisation must only use software that provides appropriate security features so it does not endanger its goals.

Nature of the Problem

Organisations are exposed to threats from actors who attack buildings, systems and people. Attacks can be personal or physical, but increasingly try to exploit vulnerabilities in IT systems. They can be direct to external interfaces (login or denial of service) or indirect (malware in e-mails or URLs).

Cradle is a part of many organisations' IT. It is helpful to discuss how Cradle's construction, architecture and features minimise the vulnerabilities that it creates for those who use it. This includes both the software itself and its use by users.

Solution

Information Assurance (IA) identifies threats, the types of actor and attack, the business impact on the confidentiality (C), integrity (I) and availability (A) of information if an attack succeeds, and defines a risk management plan to manage, mitigate or eliminate the risks from the threats. In a plan, information security (InfoSec) principles ensure that only authorised users (C) access to accurate and complete information (I) when required (A).

Benefits

Cradle's architecture and features are parts of an organisation's InfoSec approach.

Certification of 3SL 1

3SL is certified under Cyber Essentials, a UK Government scheme to confirm the basic integrity of the IT systems of those companies who interact with the UK Government, and the IASME Governance standard.

Known Origin 2

All software within, or used to build, the Cradle product has a known origin, as one of:

1. Software written by 3SL in the United Kingdom
2. Open source tools (such as compilers) for which 3SL has the source code and the rights to use to develop commercial products
3. Open source libraries for which 3SL has source code and rights to redistribute in source or compiled form
4. Commercial libraries for which 3SL has source code and rights to redistribute in source or compiled form
5. Commercial runtime libraries which are provided by 3SL or can be downloaded from their authors (such as Microsoft®)

Guaranteed Content 3

3SL can guarantee that it is not possible for malicious or otherwise unknown software to be built into the Cradle product distribution without the knowledge of 3SL.

3SL guarantees that we have not included in Cradle any software that:

- Monitors your activity for use by 3SL or any third party
- Collects your data for use by 3SL or any third party
- Sends your activity or data to 3SL or any third party
- Acts as a backdoor into your databases

Cradle uses the facilities of the host operating system (Windows® and Linux®) and other services (such as LDAP and Microsoft .NET®) only by calling publicly documented interfaces.

Every file in the Cradle distribution is virus-checked as part of the 3SL release process for every Cradle distribution. All Cradle components are white-listed with anti-virus (AV) vendors prior to release.

3SL can therefore guarantee that the Cradle product does not contain any malicious or otherwise unknown code.

Software Certification and Integrity 4

All Cradle executables on Windows are digitally signed using a certificate issued by Symantec. Such certificates confirm that 3SL is the author of each executable, and also that each executable has not been changed in any way since it was built by 3SL.

We also provide both MD5 and SHA512 checksums for the

Cradle distribution files. These can be used to confirm the integrity of Cradle software downloads, ensuring that a Cradle software distribution has not been tampered with, damaged or modified in any way in whatever method has been used to copy that file from 3SL to you.

Certification of Cradle 5

Cradle has been approved for use within the UK Government and has received security accreditation by passing an IT Health Check (ITHC) as defined by CESG and successfully completing an IA risk assessment and InfoSec verification.

The Cradle ITHC confirms the correct implementation of security functionality and confirms that there are no vulnerabilities in the web-based and non-web-based Cradle components that could compromise the C, I or A of information stored within Cradle.

Anti-Virus Software and Whitelists 6

A lot of malicious software circulates in the Internet, in malicious websites, hidden inside compromised websites unknown to their owners, attached to e-mails, or embedded in a variety of types of document (Word, Excel, PDF and others) that are attached to e-mails.

In response to this threat, there are many anti-virus (AV) products from many vendors that will scan e-mails, data exchanges with websites, files on disk, and executing images in computer memory.

It is inevitable that AV products will sometimes incorrectly report a threat. Such occurrences are termed ***false positives***, where the AV product incorrectly reports that a file (typically but not necessarily an executable) contains a virus or other malicious code, when no such threat exists. To minimise false positives, all reputable AV vendors maintain ***whitelists*** of files that they have verified. These whitelists are part of the regular updates to their products' virus definition lists.

3SL submits all the Cradle executables to the vendors of over 20 AV products, so that they can be included in their AV product whitelists. As such, AV products should never report any threat from any part of any Cradle software distribution, firstly because there are no threats in any part of Cradle, and secondly because they have verified Cradle for themselves and added Cradle's files to their AV products' whitelists.

No External Links or References 7

Cradle has no interface to any external resource, either on the Internet or otherwise. This means that:

1. Everything used by Cradle is in the local installation:

a) Software executables, libraries and related scripts
b) Documentation and on-line help files
c) Example data files
d) Utilities
e) Images, used in the Cradle UI or by users as clipart
f) Message catalogues for all languages supported in Cradle, currently Chinese, Dutch, English, French, German, Korean, Russian and Welsh

2. Cradle does not access the Internet for any purpose, including:
   a) Seeking updates to itself
   b) Verifying your licences
   c) Registering itself with 3SL or anyone else
   d) Reporting your use of Cradle

Information at Rest8

All Cradle databases are managed centrally by the Cradle Database Server (CDS) and are stored on the server hosting the CDS or in network attached storage (NAS). All information created in Cradle projects is stored in these databases. All Cradle configuration files and logs are stored on the server hosting the CDS.

All data in Cradle databases is stored unencrypted in its native format (such as plain text in UTF-8, rich text, Word documents) so it can be searched efficiently. You can specify that data is stored compressed, offering some protection. User passwords are stored encrypted with a one-way ***hash***. 3SL recommends hardware encryption of the filesystem or storage media where appropriate.

Users do not need any access to the server that hosts the CDS. Hence:

- The only means to access data in Cradle is via Cradle
- The only means to access Cradle databases, logs and configuration files is via the server that hosts the CDS

No project information, configuration or setup data of any kind is stored on Cradle users' computers, regardless of whether these users access Cradle through a web browser or a non-web client.

The only methods by which end users can produce local copies of data stored in Cradle are:

- Export
- Publish
- Database file viewer

You can control who can perform these operations. You can control which information can be exported or published. This means that you can prevent end users producing local copies of any information held in any Cradle database.

Information in Transit 9

The ***information pathways*** in Cradle are:

1. Between a user's Cradle client or API application and the CDS
2. Between a user's web browser and the Cradle Web Server (CWS)

Pathway **1** uses an internal 3SL protocol over TCP. All data is passed through this pathway using a proprietary hash. If this pathway is open to external attack, you can tunnel it through SSL, such as through a VPN connection.

Pathway **2** uses your choice of HTTP or HTTPS. Cradle includes open source tools to create the certificates for HTTPS if you decide not to purchase certificates from a Certificate Authority (CA).

Passwords are sent encrypted through these pathways, regardless of whether the pathway is encrypted, or not.

Resistance to Attack 10

Cradle resists attack by:

1. Logging of attempted attacks to the CDS
2. Automated blocking of source IPs attempting Denial of Service (DoS) attacks
3. Automatic disabling of a user login after N (default 3) consecutive failures over any period of time
4. Automatic disconnection of any user who has been idle for a period of time that you can control
5. Independent verification (see "Certification of Cradle" on page 3) that Cradle web UIs are not vulnerable, including being resistant to:
   - Remote code execution
   - Code injection
   - Format string vulnerabilities
   - Cross Site Scripting (XSS)
   - Username enumeration
   - Verbose error messages
   - Stack trace enabled in message reporting

Resilience 11

The CDS is designed to run, unattended, 24x7. There is no need for periodic downtime, even during backups. The CDS is designed to detect and recover from any data or internal error so as to be available at all times.

User Authentication 12

Each Cradle database has its own set of ***user profiles*** (login accounts) in a ***user register***. Each user is identified by a ***username*** protected by a ***password***.

Usernames identify people and their role. Usernames are used to:

- Specify who owns information, part of access control
- Record who has created or changed information
- Record who has created or changed links between information
- Record who has performed operations on information
- Define access and operation control policies

Passwords control all types of access to Cradle, which are:

- From Cradle clients
- From web browsers using a Cradle web UI
- From an application created with the Cradle API

You can enforce a range of policies for passwords:

- Minimum length
- Required characters
- Matching a pattern specified in a regular expression
- Password aging with optional warnings
- Password history with a user-defined cycle length
- Only able to change password once per day
- Must change password at next login
- Consecutive login failures disables the user profile

Cradle can reference an LDAP directory (such as FreeIPA, Microsoft's Active Directory® or OpenLDAP) to verify usernames, or username and password authentication, including ***single-sign-on***.

Access Control 13

Cradle provides access controls for:

- Information
- Operations on information

You create these controls using user-defined ***privileges*** and ***skills*** that you grant to users' user profiles and the information structure that you choose for each database (termed a ***schema***).

Information controls include:

1. Can a user browse the database, or only access it by a ***phase hierarchy*** of predefined operations that you have defined
2. Can a user access each type of information, or not
3. If a user can access a type of information, can the user access an item of that type: read-only, read-write or no access
4. If a user can access an item, can the user access each attribute in the item: read-only, read-write or no access
5. Can a user create, modify or delete an item's links

Operation controls include whether a user can or cannot:

- Define searches for information or only use searches that have been provided by you. Your searches may

not find all items that exist or that could be accessed by the user, if you wish.
- Define views of information or only use views that you have provided. Your views may not show all attributes of the items that could be accessed by the user.
- View the database files
- Export information outside Cradle
- Import information into Cradle
- Publish views of information to external files
- Publish reports of information to external files
- Publish documents of information to external files
- Create or modify analysis, architecture or design models
- Perform performance assessments of models
- Run metrics
- Produce KPI (key performance indicator) dashboards

Data Quality 14

Cradle provides several mechanisms to help ensure the quality of your project information, including:

1. ***Rule sets*** that, for example, set the value of attributes automatically based on value(s) of other attribute(s)
2. ***Calculations*** that set attributes automatically based on the attributes in the same or other linked items
3. User-defined lists of allowed values for attributes
4. ***Conformance*** checks on the quality of text statements, based on either user-defined lists of terms that improve or impair quality and/or user-defined formats to which textual statements should conform
5. User-defined quality checks for any attribute
6. ***Consistency*** check the contents of architecture, analysis or design models in MBSE (model based systems engineering) processes to ensure consistency of data, behaviour and control. Such checks can span models, for example when a model of an equipment model is referenced from an architecture.

Cradle also provides a range of ***database integrity checks*** and ***cross reference integrity checks*** that can be run to verify the integrity of every attribute of every item, and the integrity of every cross reference between every pair of items in the database.

Baselines 15

Cradle has a Configuration Management System (CMS) that allows the creation of ***baselines***. Items of information enter a baseline by a formal review process by following a user-defined ***workflow***.

All stages in these workflows are recorded in the CMS's ***configuration log*** that can be queried and reported in your quality review and governance processes. This gives

assurance that your quality processes are being executed correctly and consistently. You can design your reviews to eliminate malicious review decisions by personnel who have been compromised by external actors.

Once baselined, information is protected against change and can be used for formal project documentation.

You can formalise all changes to baselined information through ***Change Requests*** and ***Change Tasks***. These have user-defined workflows. Information changed in this way follows a series of reviews in a user-defined workflow to become part of a later baseline.

Collectively these mechanisms provide assurance that the information in your project baselines has been properly reviewed and is free from unwanted external influences.

Logging of Activity16

Cradle can record up to 1 billion changes to each item in the database. For each change, it records:

- The date and time
- The user who made the change
- Why the change occurred
- Which attributes were changed
- Old and new values of each changed attribute

Changes to items can be notified automatically by ***alerts***. Changes can be reversed, individually or in a sequence, which will add further entries to the change history.

Notifications 17

Cradle ***alerts*** are notifications of events in the database, delivered through Cradle, e-mail (IMAP and SMTP) or both. E-mail templates are provided to customise notifications sent by e-mail. You can control which events will produce an alert. These events include:

- Creation, deletion or change of specific item type(s)
- Creation, deletion or change of links between items
- Changing a category value
- Use of specific Cradle tools and operations

When combined with change logs, see “Logging of Activity” on page 7, these features allow recovery from any database activity that is authenticated and access controlled, but that is nonetheless malicious, For example, activity by personnel who have been compromised by an external threat actor.

API Misuse 18

Cradle provides an API (Application Programming Interface) through which you can create your own custom applications and tools that can operate on the contents of Cradle databases.

The API provides the same opportunities to manipulate information in a database as Cradle tools supplied by 3SL, but it also imposes the same controls and restrictions. Therefore:

- API applications login to a database with a username and password
- All authentication rules also apply to API applications
- All access controls also apply to API applications
- It is not possible for an API application to perform more malicious operations in a database than an interactive user

Therefore, the opportunities for misuse of an application using the API are no greater than the opportunities for misuse of a Cradle tool from 3SL. Such opportunities can, and should be, constrained by the authentication rules and access controls that you apply to your databases.

Problem Diagnosis19

Cradle includes a feature to ***sanitise*** exported information in which all meaningful information is replaced by x characters in the export file. This allows you to pass your data to 3SL for problem diagnosis without disclosing sensitive information. For 3SL, it is only the structure of your data and the schema that are important. This simple mechanism allows 3SL to solve customer problems when other approaches have been insufficient.