



**HYPERLEDGER**

Linux Foundation  
2021 HyperLedger Fabric  
Penetration Test

Tevora Threat Research Group

Delivered March 22, 2021

**CONFIDENTIAL**

This report is confidential for the sole use of the intended recipient(s). If you are not the intended recipient, please do not use, disclose, or distribute.

# Table of Contents

EXECUTIVE SUMMARY .....	3
PURPOSE.....	3
PROJECT PARTICIPATION .....	3
SCOPE.....	4
FINDINGS OVERVIEW.....	5
APPLICATION PENETRATION TEST RESULTS.....	5
ATTACK MAP.....	7
TECHNICAL SUMMARY .....	13
SUMMARY OF FINDINGS.....	13
TECHNICAL FINDINGS.....	14
APP-01296 GO PROFILING ENDPOINT EXPOSED .....	14
APP-01302 DEPENDENCY CONFUSION .....	18
APP-01310 FABRIC CA BOOTSTRAPPED WITH ADMIN PASSWORD IN CLEARTEXT ON COMMANDLINE .....	21
APPENDIX A: ABOUT TEVORA .....	23
APPENDIX B: SCORING OF FINDINGS.....	24
APPENDIX C: PENETRATION TESTING METHODOLOGY .....	27

# Executive Summary

## Purpose

The 2021 HyperLedger Fabric Penetration Test for Linux Foundation was conducted from 11 January, 2021 to 01 March, 2021 to help ensure Linux Foundation resources are secure from advanced threat actors.

Additional objectives for this penetration test were based on industry standard guidelines as follows:

- Identification of vulnerabilities so that they can be remediated prior to being exploited by an attacker
- Direct observation of restricted services or data in the absence of expected access controls
- Compromise of an intermediary device used by privileged users to access secure network zones
- Compromise of the domain used by privileged users
- Sensitive data leakage or exfiltration
- Verification of application logic, session handling, and API security for applications using supplied credentials
- Verification that only authorized services are exposed to the network perimeter
- Verification of network segmentation of non-privileged and privileged networks

## Project Participation

The following Linux Foundation and Tevora personnel were directly involved in this assessment.

**Sponsors** David Huseby, Linux Foundation

**Participants** Ry Jones, Linux Foundation  
Kevin Dick, Tevora

## Scope

This report contains the summary of project scope, findings, and recommendations resulting from the Application Penetration Test conducted by Tevora against the Linux Foundation environment.

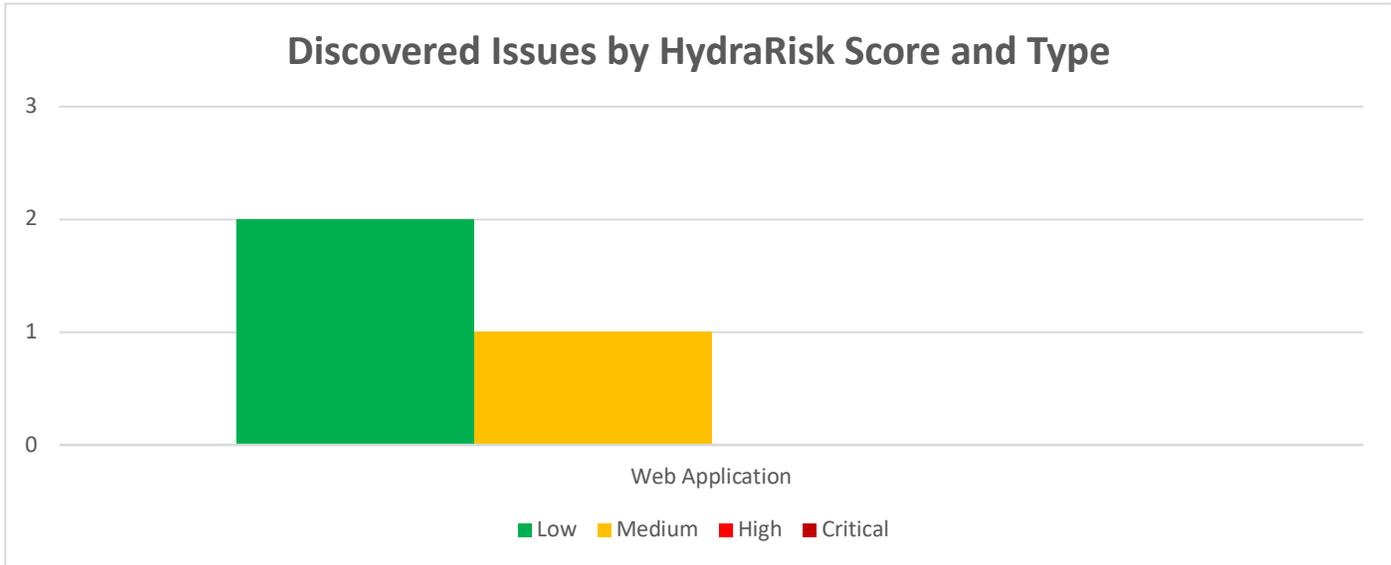
### Application Penetration Test

The following items were considered in scope:

#### **HyperLedger Fabric:**

- <https://github.com/hyperledger/fabric>

# Findings Overview



## Application Penetration Test Results

Tevora performed penetration testing of the Hyperledger Fabric blockchain platform, analyzing the security of Peers, Orderers, Fabric CA, and their intercommunications. Dynamic and static analysis was performed including debugging of critical security code paths such as identity checking and endorsement validation, as well as testing live test networks' resilience against potential attack scenarios. Tevora discovered some issues during testing, most notably the exposure of Go's PPROF profiling endpoint by Peer and Orderer services in the sample test network configurations, however, no serious issues were found with the core security architecture of the platform nor implementation of security-critical components.

Tevora discovered Go's PPROF endpoint to be enabled in the test-network for Peers and Orderer nodes, which, if the test-network is used as a base configuration for production deployments, could lead to this endpoint being exposed in production. If exposed, this debug endpoint potentially exposes the affected peer and orderer components to denial of service or informational leakage risks. Hyperledger has modified the sample configurations to disable this service by default.

Tevora focused much of the initial testing on validating Fabric's trust model with the default MSP, ensuring that authentication and authorization checks were enforced as expected. Tevora forged certificates and sent a variety of malformed requests to evaluate the resilience of authentication and authorization flows to attack. Tevora found that, in the default MSP, validation was performed as expected and was resilient to forgery attempts. This may be due in part to Fabric's use of standard mechanisms such as .x509 certificates and proven public/private trust to authenticate components. Tevora found the default MSP to be sound, and the flexibility in MSP implementation welcome.

Additionally, a number of potential attack scenarios were explored including compromised peers, transaction replay, unauthenticated denial of service, gossip replication exploitation, unauthorized chaincode installation, and endorsement spoofing. Tevora performed both experimental analysis and direct debugging of security critical code paths. Tevora performed experiments using modified peers and clients to launch attack scenarios then subsequently used instrumentation such as the Hyperledger Blockchain Explorer to analyze the results of the attempted attacks. Tevora also stepped through security-critical code paths in a debugger to analyze directly how validation checks are performed. In all cases, Tevora did not observe any deviation from expected behavior, nor were any attack paths apparent after direct analysis of the execution flow. Lastly, Tevora extended testing significantly beyond the planned timeframe to ensure good coverage was performed for the scenarios and a good understanding of the validation behavior was achieved.

Tevora performed enumeration of and dynamic testing against exposed endpoints on Fabric components, including gRPC endpoints, Fabric CA REST APIs, and Peer health web endpoints. Tevora observed which services are accessible in default configurations and performed penetration testing against them. The majority of complexity was, as expected, in the gRPC endpoints, which were thoroughly tested. Other endpoints were found to present a minimal attack surface. Arguably the most critical of the discovered web services endpoints is the Fabric CA REST APIs. These were found to be resilient to attack and observed to implement recommended hardening such as restriction of failed login attempts to prevent password brute-forcing. No serious issues were observed here, and many of these endpoints can be restricted in production deployments, leaving only gRPC communications open to other organizations.

Tevora discovered Go's PPROF endpoint to be opened on 0.0.0.0:6060 by default for Peers and Orderer nodes, potentially exposing these components to denial of service or informational leakage risks. In many cases, production deployments will block access to these ports at the network layer; however, the service being exposed by default increases the risk of insecure deployments. In general, Tevora recommends hardening default configurations, with debug functionality being opt-in.

During testing, Tevora was alerted to a dependency confusion issue in the fabric-chaincode-node SDK reported via Hyperledger's bug bounty program. Tevora validated this issue was present, and performed a sweep of that node SDK, and other Fabric related SDKs for this or other issues. In-depth testing on all SDKs was not performed to the degree of the core fabric repo. However, Tevora found few issues when sweeping the SDKs.

Overall, Tevora found Fabric to be a robust, well-architected, and similarly well-implemented platform. It is apparent that Fabric has reached a mature state of security and functionality and is well-positioned for updates should any flaws be discovered in the underlying industry-standard cryptography used for the default MSP and validation, partly thanks to Fabric's modular and open-source nature. The issues discovered by Tevora all relate to increasing hardening of default configurations to avoid insecure deployments of Fabric in the field. In general, Tevora anticipates the highest likelihood of attack targets to be implementation errors such as weak passwords, exposed debug interfaces, or trust misconfigurations. The Hyperledger Fabric team may want to add security guidelines and considerations into the "deploying a production network" documentation, because the example

networks, although quite well documented and complete, do potentially dangerous things like setting weak administrative passwords.

## Attack Map

Tevora analyzed the footprint and expected deployments of the fabric platform to develop potential attack scenarios and test cases to evaluate these scenarios. Overall, Tevora found Fabric to be resilient to test cases, with no major application logic or design flaws discovered outside of legacy (1.4) chaincode installation processes. Tevora performed many application logic tests on the fabric platform and has highlighted the attack scenarios below that Tevora believes to be most critical to validating Fabric's security model. In all cases, Tevora found Fabric to be resilient to attack and spoofing attempts.

### Identify Spoofing

A key tentpole of Fabric is trust, with all subsequent security requiring the trust mechanisms to be correctly implemented. Unlike public blockchain systems using proof of work to establish trust, Fabric relies on being able to identify and authenticate organizations and their components. Because of this, a large amount of testing was focused on fabric Membership Service Providers (MSPs) and how components authenticate to each other.

Tevora analyzed how peers communicate and authenticate with certificates, stepping through the default certificate-based MSP provider to ensure identification and validation of clients is done correctly. Tevora both created spoofed certificate test cases and stepped through the MSP validation code in a debugger to analyze behavior for areas of potential weakness. Tevora used modified peers to both test how data is handled by interesting functions directly and to modify Peer CLI commands to send test cases to Peer servers.

#### ***Example snippet of modified peer used as a scratchpad to test the handling of unexpected cert data in various stages of the validation pipeline***

```

203     return nil, errors.New(message, "encoding of identity failed")
204 }
205 // We serialize identities by prepending the MSPID and appending the ASN.1 DER content
206
207 //going to just replace the pem bytes with public key of the admin
208 pubkey_to_spoof_path := "/home/kdick
209 admin_cert, err := readFile(pubkey_to_spoof_path)
210
211
212
213 sId := &msp.SerializedIdentity{Mspid: id.id.Mspid, IdBytes: pemBytes}
214
215 //override with a nice evil cert
216 sId = &msp.SerializedIdentity{Mspid: id.id.Mspid, IdBytes: admin_cert}
217
218
219
220 //tevora hackery
221 //sId = &msp.SerializedIdentity{Mspid: [1,2,3] , IdBytes: pemBytes}
222
223
224 idBytes, err := proto.Marshal(sId)

```



Tevora noted that some caching is performed of MSP cert validation, presumably to optimize performance. Tevora evaluated this to ensure a spoofed certificate could not be sent after an initial cache but found that signatures are always validated. The cache key is a SHA digest of the entire raw public cert; any modification to the certificate will cause a cache invalidation, and the certificate is always used to validate the signature. There were no issues identified with caching.

```

166 if up.SignedProposal.Signature == nil { errors.Errorf("empty signature bytes") }
169
170 // get the identity of the creator
171 creator, err := idDeserializer.DeserializeIdentity(up.SignatureHeader.Creator)
172 if err != nil {
173     logger.Warningf(template: "access denied: channel %s", err)
174     return errors.Errorf("access denied: channel [%d] creator org unknown, creator is malformed", up.ChannelID())
175 }
176
177 genericAuthError := errors.Errorf("access denied: channel [%d] creator org [%s]", up.ChannelID(), creator.GetMSPIdentifier())
178 // ensure that creator is a valid certificate
179 err = creator.Validate()
180 if err != nil {
181     logger.Warningf(template: "access denied: identity is not valid: %s", err)
182     return genericAuthError
183 }
184
185 logger = logger.With(args... "mspID", creator.GetMSPIdentifier())
186
187 logger.Debug(args... "creator is valid")
188
189 // validate the signature
190 err = creator.Verify(up.SignedProposal.ProposalBytes, up.SignedProposal.Signature)
191 if err != nil {
192     logger.Warningf(template: "access denied: creator's signature over the proposal is not valid: %s", err)
193     return genericAuthError
194 }
195
196 logger.Debug(args... "signature is valid")
    
```

No other notable leads were identified while stepping through the code flow, and the validation process remained resilient to sanity test cases and experiments with forged certificates, modified proposals, and modified headers. The use of industry-standard public/private key cryptography as well as a concise and effective validation flow benefit Fabric and provide no observed avenues of attack other than potential vulnerabilities in the underlying cryptography used.

In the case of RSA, ECDSA, or ED25519 implementations being broken (or underlying algorithm being broken i.e. in RSA's case an algorithm for fast factoring primes), Hyperledger would be far from the only platform to be affected by this, and new solutions would likely be easy to plug in, most likely with a simple update to Go's x509 library. Similarly, vulnerable methods can be deprecated and made unsupported.

Additionally, Tevora observed a flexible MSP interface that allows other implementations to be developed relatively seamlessly. Although this could lead to untested and unsecured deployments with vulnerable MSP

implementations, the most Hyperledger may want to do here is warn users of the diligence required when rolling their own MSP.

Lastly, Fabric also makes optional use of TLS wrappers for connections between components. This extra layer of security encrypts gRPC traffic in transit and mitigates the risk of any potential future man in the middle (MiTM) attacks against the MSP itself since the MSP communication occurs within the protection of the TLS tunnel.

## Malicious ChainCode Install

Because chaincode, by design, executes defined code on targeted peers (typically in a Docker container), Tevora wanted to ensure that unauthenticated clients or malicious evil peers in a network could not deploy malicious chaincode to the network. Any vectors for code execution are a valuable target for attackers, so it is important that these chaincodes are only deployable with proper administrator authorization.

The 1.4 chaincode lifecycle does allow peers from other organizations to replicate potentially malicious chaincode across org boundaries. Other security measures, such as isolation of executed chaincode to Docker containers, are in place; however, this could lead to denial of service or network-based attacks from the compromised container. Fabric is already aware of these risks, as evidenced by documentation and public presentations (<https://www.youtube.com/watch?v=XvEMDScFU2M&t=1407s>

). Tevora recommends migrating to 2.0 chaincode lifecycles for existing implementations.

The 2.0 model of chaincode installation was found to be designed with both security and scalability in mind, requiring organization admins to install specific code bundles to peers and not replicating the code packages across peers. In addition to the major security improvements, 2.0 chaincode lifecycle allows different chaincode to be implemented for the same endorsed interface greatly increasing the flexibility of chaincode operations by allowing per organization customizations. Tevora did not identify any ways to bypass authentication to install chaincode to peers without having the requisite admin access, nor was chaincode observed to be replicated as in the 1.4 protocol.

## Endorsement Spoofing

Tevora validated Fabric's distributed endorsement validation architecture by simulating malicious peers and client applications modified to attempt to commit and replicate transactions that have failed the approved endorsement policy. Tevora used the default majority endorsement policy for these tests, modifying evil peers and clients to propose and commit to the ledger transactions not meeting the endorsement policy.

Tevora observed that the unmodified peers successfully enforced their approved endorsement policy on chaincode execution and replication.

## Example attempted \_lifecycle chaincode endorsement replication attempt result

Org2MSP	mychannel	8a218c...	ENDORSER_TRANSACTION	papercontract	2021-02-26T03:15:36.040Z
Org2MSP					02-26T03:15:32.800Z
Org2MSP					02-26T03:14:30.893Z
Org2MSP					02-26T03:13:10.270Z
Org2MSP					02-26T03:12:57.109Z
Org2MSP					02-26T03:12:42.842Z
Org2MSP					02-26T02:09:56.211Z
Org2MSP					02-26T02:09:39.795Z
Org2MSP					02-26T02:05:41.145Z
Org2MSP	mvchannel	daa9cd...	ENDORSER_TRANSACTION	lifecvcl	2021-02-16T03:41:01.380Z

**Transaction Details** ✖

<b>Transaction ID:</b>	cbed5e26c92082cee9d3ec929b5093e5b6020736bb18e35ddebe13456e3eafb7 <span style="float: right;">🔗</span>
<b>Validation Code:</b>	ENDORSEMENT_POLICY_FAILURE
<b>Payload Proposal Hash:</b>	00b3000957fa89c59b1360c9cc1936bd4889f13d32af98093f11a195d1349d10
<b>Creator MSP:</b>	Org2MSP
<b>Endorser:</b>	{"Org2MSP"}
<b>Chaincode Name:</b>	_lifecycle
<b>Type:</b>	ENDORSER_TRANSACTION
<b>Time:</b>	2021-02-26T02:09:56.211Z
<b>Reads:</b>	<ul style="list-style-type: none"> <li>▼ root: [] 1 item</li> <li>▼ 0: {} 2 keys                             <ul style="list-style-type: none"> <li>chaincode: "_lifecycle"</li> <li>▶ set: [] 5 items</li> </ul> </li> </ul>
<b>Writes:</b>	<ul style="list-style-type: none"> <li>▼ root: [] 1 item</li> <li>▶ 0: {} 2 keys</li> </ul>

## Same transaction from "malicious" Peer's ledger

Org2MSP	mychannel	4d5d0c...	ENDORSER_TRANSACTION	papercontract	2021-02-26T03:14:30.893Z
Org2MSP					02-26T03:13:10.270Z
Org2MSP					02-26T03:12:57.109Z
Org2MSP					02-26T03:12:42.842Z
Org2MSP					02-26T02:09:56.211Z
Org2MSP					02-26T02:09:39.795Z
Org2MSP					02-26T02:05:41.145Z
Org2MSP					02-16T03:41:01.380Z
Org2MSP					02-16T03:40:19.469Z

**Transaction Details** ✖

<b>Transaction ID:</b>	cbed5e26c92082cee9d3ec929b5093e5b6020736bb18e35ddebe13456e3eafb7 <span style="float: right;">🔗</span>
<b>Validation Code:</b>	VALID
<b>Payload Proposal Hash:</b>	00b3000957fa89c59b1360c9cc1936bd4889f13d32af98093f11a195d1349d10
<b>Creator MSP:</b>	Org2MSP
<b>Endorser:</b>	{"Org2MSP"}
<b>Chaincode Name:</b>	_lifecycle
<b>Type:</b>	ENDORSER_TRANSACTION
<b>Time:</b>	2021-02-26T02:09:56.211Z
<b>Reads:</b>	<ul style="list-style-type: none"> <li>▼ root: [] 1 item</li> <li>▶ 0: {} 2 keys</li> </ul>
<b>Writes:</b>	<ul style="list-style-type: none"> <li>▼ root: [] 1 item</li> <li>▶ 0: {} 2 keys</li> </ul>

## Transaction Replay

Tevara reviewed the Peer endorsement lifecycle architecture, observed transaction processing with Transaction ID validation, and tested with repeated messages to validate the expected behavior of a Fabric deployment. As designed, duplicate transactions cannot be currently replayed in the tested Fabric 2.3 test-network.

### Analyzing TXID (Transaction ID) generation and validation

Tevara found that TXIDs are made using a client-generated nonce, which is concatenated with the creator (essentially public key for the authenticated client) and hashed. Both the nonce and creator info is sent in the authenticated payload contents to Peers, allowing peers to confirm the correct transaction ID.

```

135 // In fact, anything coming to the Propose API is by definition an endorser
136 // transaction, so any other header type seems like it ought to be an error... oh well.
137
138 default:
139     return errors.Errorf("invalid header type %s", common.HeaderType(up.ChannelHeader.Type))
140 }
141
142 // ensure the epoch is 0
143 if up.ChannelHeader.Epoch != 0 { errors.Errorf("epoch is non-zero")}
144
145 // ensure that there is a nonce
146 if len(up.SignatureHeader.Nonce) == 0 { errors.Errorf("nonce is empty")}
147
148 // ensure that there is a creator
149 if len(up.SignatureHeader.Creator) == 0 { errors.Errorf("creator is empty")}
150
151 expectedTxID := protoutil.ComputeTxID(up.SignatureHeader.Nonce, up.SignatureHeader.Creator)
152 if up.TxID() != expectedTxID {
153     return errors.Errorf("incorrectly computed txid '%s' -- expected '%s'", up.TxID(), expectedTxID)
154 }
155
156 if up.SignedProposal.ProposalBytes == nil { errors.Errorf("empty proposal bytes")}
157
158 if up.SignedProposal.Signature == nil { errors.Errorf("empty signature bytes")}
159
160 // get the identity of the creator
161 creator, err := idDeserializer.DeserializeIdentity(up.SignatureHeader.Creator)
162 if err != nil {
163     logger.Warningf("access denied: channel %s", err)
164     return errors.Errorf("access denied: channel [%s] creator org unknown, creator is malformed")
165 }
166
167 genericAuthError := errors.Errorf("access denied: channel [%s] creator org [%s]", up.ChannelID(), creator.GetMSPIdentifier())
168 // ensure that creator is a valid certificate
169 err = creator.Validate()
170 if err != nil {
171     logger.Warningf("access denied: identity is not valid: %s", err)
172 }
173 *UnpackedProposal.Validate(idDeserializer, msp.IdentityDeserializer)
174
175 // ComputeTxID computes TXID as the hash computed
176 // over the concatenation of nonce and creator.
177 func ComputeTxID(nonce, creator []byte) string {
178     // TODO: Get the Hash function to be used from
179     // channel configuration
180     hasher := sha256.New()
181     hasher.Write(nonce)
182     hasher.Write(creator)
183     return hex.EncodeToString(hasher.Sum(b: nil))
184 }
185
186 // CheckTxID checks that txid is equal to the Hash computed
187 // over the concatenation of nonce and creator.
188 func CheckTxID(txid string, nonce, creator []byte) error {
189     computedTxID := ComputeTxID(nonce, creator)
190
191     if txid != computedTxID { errors.Errorf("invalid txid. got [%s], expected [%s]", txid, computedTxID) }
192
193     return nil
194 }
195
196 // Here we handle uniqueness check and ACLs for proposals targeting a chain
197 // Notice that ValidateProposalMessage has already verified that TXID is computed properly
198 if _, err = e.Support.GetTransactionByID(up.ChannelHeader.ChannelID, up.ChannelHeader.TxID); err == nil {
199     // increment failure due to duplicate transactions. Useful for catching replay attacks in
200     // addition to benign retries
201     e.Metrics.DuplicateTxFailure.With(meterLabels...).Add(delta 1)
202     return errors.Errorf("duplicate transaction found [%s]. Creator [%s]", up.ChannelHeader.TxID, creator)
203 }

```

### Validated duplicate check in core/endorser/endorser.go

```

202 // Here we handle uniqueness check and ACLs for proposals targeting a chain
203 // Notice that ValidateProposalMessage has already verified that TXID is computed properly
204 if _, err = e.Support.GetTransactionByID(up.ChannelHeader.ChannelID, up.ChannelHeader.TxID); err == nil {
205     // increment failure due to duplicate transactions. Useful for catching replay attacks in
206     // addition to benign retries
207     e.Metrics.DuplicateTxFailure.With(meterLabels...).Add(delta 1)
208     return errors.Errorf("duplicate transaction found [%s]. Creator [%s]", up.ChannelHeader.TxID, creator)
209 }

```

# Technical Summary

## Summary of Findings

**Total Penetration Test Findings 3**

<b>Application Findings</b>	<b>Status</b>	<b>HydraRisk</b>	
APP-01302 Dependency Confusion	<b>Remediated</b>	<b>12</b>	<b>Medium</b>
APP-01296 Go Profiling Endpoint Exposed	<b>Remediated</b>	<b>10</b>	<b>Low</b>
APP-01310 Fabric CA Bootstrapped with Admin Password in Cleartext on Commandline	Discovered	<b>6</b>	<b>Low</b>

# Technical Findings

## APP-01296 Go Profiling Endpoint Exposed

### Description

While reviewing Fabric's source files, Tevora discovered imports of "net/http/pprof". When this item is imported, it can open a debug server on port 6060, which an attacker can use to enumerate:

- Function names and file paths
- Profiling data
- Denial of Service risk via profiling load

Tevora observed this debug server on the test-network instances, though not directly accessible unless port forwarding rules were added. The Fabric team has updated the test-network configuration to no longer enable this debug interface.

A hyperledger fabric administrator may inadvertently leave this port exposed in production if they are not familiar with the risks of this interface.

<b>Status</b>	Discovered	<b>CVSS Base Score</b>	<b>7.3 High</b>	<b>HydraRisk</b>	<b>13 Medium</b>
				Consequence	4
				Probability	3
				Velocity	2
				Criticality	3
				Responsiveness	1

### Affected Systems

The following systems have been identified and are affected:

Repo	File
https://github.com/hyperledger/fabric-ca	fabric-ca/lib/server.go
https://github.com/hyperledger/fabric	fabric/cmd/peer/main.go
https://github.com/hyperledger/fabric	fabric/orderer/common/server/main.go



## Viewing PPROF listener on Org1 sample peer container

```

-> ~ docker ps -a
CONTAINER ID   IMAGE                                COMMAND                  CREATED        STATUS
90b497813bcc   gliderlabs/logspout                "/bin/logspout"         2 days ago    Up 2 days
4b7575cce14d   hyperledger/fabric-peer:latest     "peer node start"      10 days ago   Exited (0) 8 days a
c88665be662d   hyperledger/fabric-peer:latest     "peer node start"      10 days ago   Up 2 days
db20bf31aae9   couchdb:3.1.1                      "tini -- /docker-ent..." 10 days ago   Up 2 days
40078c7fd61c   hyperledger/fabric-orderer:latest  "orderer"               10 days ago   Up 2 days
821eb589c05e   couchdb:3.1.1                      "tini -- /docker-ent..." 10 days ago   Up 2 days
999372985fc4   hyperledger/fabric-ca:latest       "sh -c 'fabric-ca-se..." 10 days ago   Up 2 days
e9c287d84090   hyperledger/fabric-ca:latest       "sh -c 'fabric-ca-se..." 10 days ago   Up 2 days
e13bed5be281   hyperledger/fabric-ca:latest       "sh -c 'fabric-ca-se..." 10 days ago   Up 2 days
29514b611435   portainer/portainer-ce             "/portainer"            3 weeks ago   Up 2 days

To get more help with docker, check out our guides at https://docs.docker.com/go/guides/

-> ~ docker exec -it c88665be662d /bin/sh
/opt/gopath/src/github.com/hyperledger/fabric/peer # netstat -plnt
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State       PID/Program name
tcp    0      0 127.0.0.1:35981        0.0.0.0:*               LISTEN     -
tcp    0      0 127.0.0.1:9443        0.0.0.0:*               LISTEN     1/peer
tcp    0      0 :::6060                :::*                   LISTEN     1/peer
tcp    0      0 :::7052                :::*                   LISTEN     1/peer
tcp    0      0 :::7051                :::*                   LISTEN     1/peer
    
```

## Accessing PPROF debug interface on local (non Dockerized) peer from another system on the same lan

Types of profiles available:

- Count Profile
- 45 [allocs](#)
- 0 [block](#)
- 0 [cmdline](#)
- 129 [goroutine](#)
- 45 [heap](#)
- 0 [mutex](#)
- 0 [profile](#)
- 12 [threadcreate](#)
- 0 [trace](#)
- [full goroutine stack dump](#)

Profile Descriptions:

- allocs:** A sampling of all past memory allocations
- block:** Stack traces that led to blocking on synchronization primitives
- cmdline:** The command line invocation of the current program
- goroutine:** Stack traces of all current goroutines
- heap:** A sampling of memory allocations of live objects. You can specify the gc GET parameter to run GC before taking the heap sample.
- mutex:** Stack traces of holders of contended mutexes
- profile:** CPU profile. You can specify the duration in the seconds GET parameter. After you get the profile file, use the go tool pprof command to investigate the profile.
- threadcreate:** Stack traces that led to the creation of new OS threads
- trace:** A trace of execution of the current program. You can specify the duration in the seconds GET parameter. After you get the trace file, use the go tool trace command to investigate the trace.

## References

- Your pprof is showing
  - <https://mmcloughlin.com/posts/your-pprof-is-showing>

- pprof exposed on kubelet healthz port
  - <https://github.com/kubernetes/kubernetes/issues/81023>

## Recommendations

In general, Tevora recommends default configurations be hardened, with debug interfaces or other potentially dangerous functionality being opt-in instead of on by default.

The PPROF server import should be removed, and PPROF loads to be done only when opted by the administrator, perhaps via an environment variable or config.yaml specification. Additionally, Tevora recommends the PPROF server be opened on localhost only by default unless otherwise specified.

# APP-01302 Dependency Confusion

## Description

During testing, another party discovered and disclosed library imports potentially affected by dependency confusion in the Node chaincode SDKs. An attacker can register the package names in public repositories, which will take precedence over the intended local package files. This could allow code execution on a system which installs the affected SDKs. Another dependency issue was also detected by Tevora; however, it is not exploitable due to NPM similar name rules preventing the public package from being registered.

<b>Status</b>	Discovered	<b>CVSS Base Score</b>	<b>7.1</b>	<b>High</b>	<b>HydraRisk</b>	<b>12</b>	<b>Medium</b>
					Consequence	4	
					Probability	2	
					Velocity	3	
					Criticality	2	
					Responsiveness	1	

## Affected Systems

The following systems have been identified and are affected:

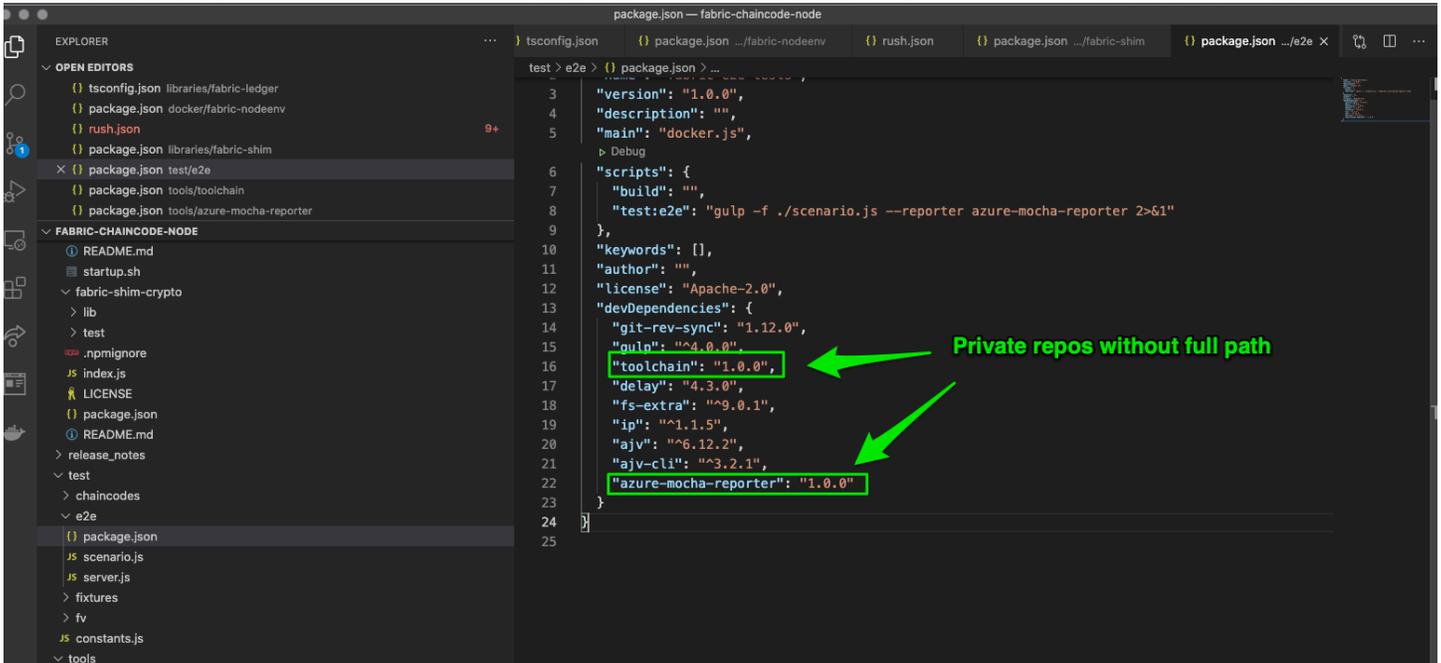
Repo	Package Type	Vulnerable Dependency	File
https://github.com/hyperledger/fabric-chaincode-node	NPM	azure-mocha-reporter	libraries/fabric-ledger/package.json libraries/fabric-shim/package.json libraries/fabric-shim-crypto/package.json test/e2epackage.json test/fv/package.json
https://github.com/hyperledger/fabric-chaincode-node	NPM	toolchain	docker/fabric-nodenv/package.json test/e2epackage.json

## Details

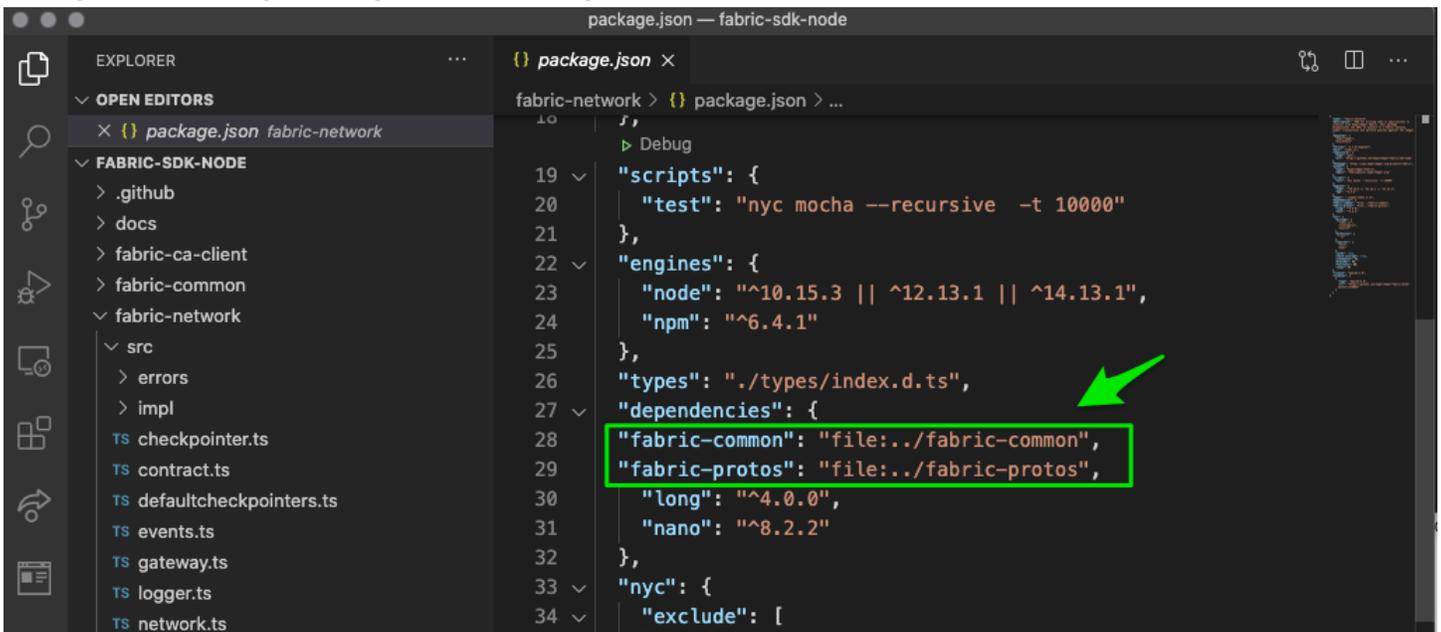
### Ticket for 'azure-mocha-reporter is vulnerable to package confusion':

<https://jira.hyperledger.org/browse/FABCN-441>

### 'toolchain' and 'azure-mocha-reporter' local repositories imported without full path in hyperledger/fabric-chaincode-node/test/e2e/package.json



### Example correct explicit import of local dependencies in fabric-sdk-node/fabric-network



## Attempting to register 'toolchain' repo and being blocked by npm similar name rules

```
npm ERR! 403 403 Forbidden - PUT https://registry.npmjs.org/toolchain - Package name too similar to existing packages; try renaming your package to '@kdic  
ktevora/toolchain' and publishing with 'npm publish --access=public' instead
```

## References

- Dependency Confusion: How I Hacked Into Apple, Microsoft and Dozens of Other Companies
  - <https://medium.com/@alex.birsan/dependency-confusion-4a5d60fec610>
- 3 Ways to Mitigate Risk When Using Private Package Feeds
  - <https://azure.microsoft.com/mediahandler/files/resourcefiles/3-ways-to-mitigate-risk-using-private-package-feeds/3%20Ways%20to%20Mitigate%20Risk%20When%20Using%20Private%20Package%20Feeds%20-%20v1.0.pdf>

# APP-01310 Fabric CA Bootstrapped with Admin Password in Cleartext on Commandline

## Description

The Fabric CA Server must be bootstrapped with an admin password if not using LDAP. Currently, this is done via the -b command-line argument providing both a username and a password as input.

Tevora advises against CLI arguments with cleartext passwords as a default bootstrapping scheme because this can lead to the password being stored in insecure locations such as git repositories or centralized audit logs. Tevora commonly observes and takes advantage of this behavior in other applications during infrastructure penetration tests.

This is not a major issue, as it typically requires multiple security failures in the deployed environment unrelated to Fabric to be exploitable in practice, however, because the Fabric CA is a tentpole of a Fabric organization's security, Tevora suggests implementing a more stringent bootstrapping procedure by default.

<b>Status</b>	Discovered	<b>CVSS Base Score</b>	3.9	Low	<b>HydraRisk</b>	6	Low
					Consequence	2	
					Probability	1	
					Velocity	1	
					Criticality	1	
					Responsiveness	1	

## Affected Systems

The following systems have been identified and are affected:

### Repo

<https://github.com/hyperledger/fabric-ca>

Details

Plaintext password default bootstrap instructions on <https://hyperledger-fabric-ca.readthedocs.io/en/release-1.4/users-guide.html#initializing-the-server>

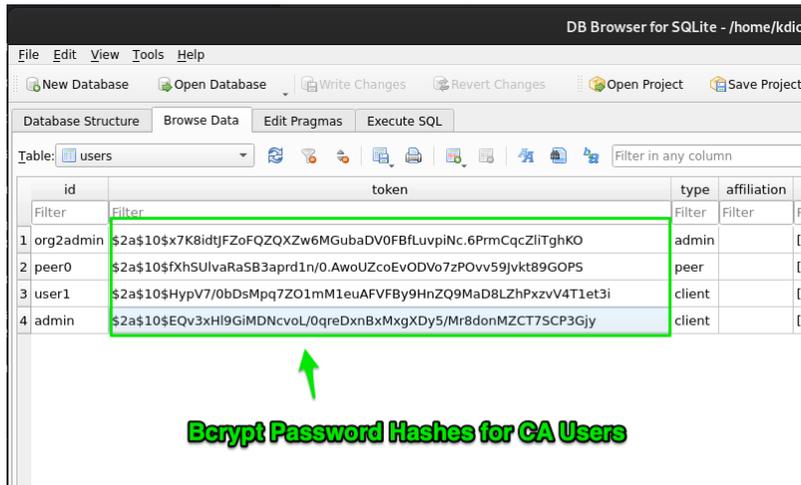
### Starting the server

Start the Fabric CA server as follows:

```
fabric-ca-server start -b <admin>:<adminpw>
```

### Passwords stored as Bcrypt in Fabric CA database

Bcrypt hashes are a great choice, and could potentially be specified instead of the cleartext in the command line parameter for additional security.



### Recommendations

Tevora recommends, by default, generating a strong random password for the CA admin account and presenting it to the user (or orchestration software) on standard out.

For pushing custom passwords directly, Tevora recommends allowing users to specify a strong hash (such as the Bcrypt observed in the actual SQLite database) to push the password, rather than a plaintext password specification.

Overall, this would only be a marginal security improvement, so Hyperledger should weigh it against any difficulties it may cause users in implementation.

## Appendix A: About Tevora

Tevora is a leading management consulting firm specializing in enterprise risk, compliance, information security solutions, and threat research. We offer a comprehensive portfolio of information security solutions and services to clients in virtually all industries and serve institutional and government clients.

Tevora's leaders are professionals with years of experience and records of accomplishments in technology as well as business. This dual background means that we understand the importance of growth and profitability and our solutions are designed to enhance both.

As a consulting firm that can fully implement whatever it recommends, Tevora works with all the industry's top vendors, yet is beholden to none. Our work and dedication have established us as a reliable partner CTOs CIOs, and CISOs can depend on to help protect against threats, both internal and external. With Tevora as a partner, business leaders can devote their energies to enhancing the overall value of information technology to their enterprise.

Tevora is a Qualified Security Assessor (QSA) and Payment Application Qualified Security Assessor (PA-QSA) in good standing with the PCI Security Standards Council. Tevora is also a DVBE (Disabled Veteran Business Enterprise) certified by the California General Services Department (Cert REF# 32786). For more information, please visit [www.tevora.com](http://www.tevora.com).

### Report Content

This report has been compiled for the exclusive use of Linux Foundation. Care has been taken to ensure that all report content and recommendations are of the highest quality and are based on sound analysis, research, and experience. Please direct any questions or concerns about the content of this report to Clayton Riness at [criness@tevora.com](mailto:criness@tevora.com).



Clayton Riness, Managing Director

# Appendix B: Scoring of Findings

Penetration Test findings are qualified using the CVSS Version 3,1 Base Score and the Tevora proprietary HydraRisk model.

## CVSSv3.1 Scoring

The CVSS version 3.1 vulnerability scoring system produces a base vulnerability score based on an Impact, and Exploitability metrics. This score is recorded for all applicable findings and is intended to provide an objective, industry-standard view of the vulnerabilities that have been found and potentially exploited.

Scoring guidelines:

- The CVSS version 3 Temporal and Environmental score metrics are not used in this report. Those factors are captured in the HydraRisk scoring model.
- In cases when multiple vulnerabilities with differing CVSS scores are summarized into a single finding, the highest contributing CVSS score is used for that finding.
- Some findings may not be given a CVSS since there is no known vulnerability but where an issue was found with the in-scope environment which differs from industry best practices or which may be used in combination with other findings to exploit a system.

## HydraRisk Scoring

Enterprise risk management is an enterprise approach to addressing the culture, processes and structures that are directed towards effective management of potential opportunities and adverse effects as they relate to risk. Taking control of informed risks allows for risks to be identified, analyzed, evaluated, treated, and monitored.

Tevora's proprietary HydraRisk Model is founded on extensive experience in enterprise risk management which has been adapted for the scoring penetration testing results. The HydraRisk score is the sum of the score for all five factors defined as follows.

**Consequence** The information security impact a threat and/or exploit has on the organization.

- 1 Trivial: Non-vital information disclosure: email addresses, WHOIS info, etc.
- 2 Reasonable: Disclosure of non-public but non-vital information
- 3 Significant: Non-privileged system access
- 4 Intolerable: Privileged system access through exploit, pivoting, or escalation
- 5 Major: Exfiltration of data: PCI, PII, intellectual property, etc.

**Probability** The likelihood of the vulnerability/threat to be exploited.

- 1 Low: No known exploit, requires skilled attacker creating a new 0-day
- 2 Unlikely: Exploit only possible using specialized tools
- 3 Moderate: Exploit is possible using common attacks or attack chaining
- 4 High: Easy to exploit by low skilled penetration tester using common tools
- 5 Critical: Easy to exploit with simple tools that are readily available

**Velocity** Assessment of how quickly a vulnerability could be exploited.

- 1 Protracted: Requires brute forcing crypto, application fuzzing, etc. over extended period
- 2 Slow: Requires extensive rainbow tables or other reference libraries to exploit
- 3 Moderate: Requires readily available reference libraries or casual observation to exploit
- 4 Quick: Requires casual observation to discover exploit
- 5 Immediate: Vulnerability can be discovered and exploited readily

**Criticality** The depth and breadth of the impact including the types of systems compromised or affected by exploiting this vulnerability.

- 1 Trivial: vulnerability affects unimportant systems: ancillary support systems
- 2 Reasonable: exploitation affects access to DMZ or other highly segmented hosts
- 3 Significant: exploitation affects access to loosely segmented hosts or client environment
- 4 Intolerable: exploitation affects substantial portions of the environment and data
- 5 Major: exploitation affects access to critical data, data integrity, and availability

**Responsiveness** The time required to treat and prevent the exploit from occurring.

- 1 Excellent: vulnerability patch or reconfiguration for exploit is readily available
- 2 Good: vulnerability patch is in development or a workaround is available
- 3 Moderate: patching, reconfiguration, and/or infrastructure re-architecting is required
- 4 Fair: infrastructure modification and/or downtime required to remediate
- 5 Poor: major infrastructure modification and/or downtime required to remediate

## Scoring Key

The following scoring key is used throughout this report, with CVSS scores ranging from 0-10 while HydraRisk scores range from 5-25.

Risk Rating	HydraRisk Score	Risk Rating	CVSS Score
<b>Critical</b>	21-25	<b>High</b>	7.0-10.0
<b>High</b>	16-20	<b>Medium</b>	4.0-6.9
<b>Medium</b>	11-15	<b>Low</b>	0.0-3.9
<b>Low</b>	5-10		

All findings are categorized as follows:

Status	Description
<b>Informational</b>	No security risk present
<b>Discovered</b>	Security risk discovered and verified, but not successfully exploited
<b>Exploited</b>	Security risk successfully exploited with proof of concept attack

## Penetration Testing Tools

Tevora employs many tools during penetration test to assist and complement manual testing including:

- Nessus Professional
- BurpSuite Pro
- ZAP (Zed Attack Proxy)
- SQLmap
- Acunetix
- NetSparker
- Custom Python scripts
- DirBuster

# Appendix C: Penetration Testing Methodology

Tevora employs a standard methodology to ensure a repeatable level of quality in all assessments. Tevora's testing methodology is based on the Penetration Testing Execution Standard (PTES)<sup>1</sup>, OWASP testing guide v4<sup>2</sup>, and years of experience in network, web, and application penetration testing.



<sup>1</sup> [http://www.pentest-standard.org/index.php/Main\\_Page](http://www.pentest-standard.org/index.php/Main_Page)

<sup>2</sup> [https://www.owasp.org/images/5/52/OWASP\\_Testing\\_Guide\\_v4.pdf](https://www.owasp.org/images/5/52/OWASP_Testing_Guide_v4.pdf)

## Phase 0: Planning and Preparation

A successful penetration test begins with planning and preparation. During this phase, Tevora works with the Client to identify the scope and any prerequisites to project execution. Tevora performs the following pre-engagement activities to prepare for testing:

- **Scope Identification:** Tevora and the Client identify the in-scope targets to be tested.
- **Testing Window Identification:** The Client provides the range of acceptable testing windows and Tevora decides when the testing will be performed within that range.
- **Objective Identification:** Tevora and the Client discuss and agree on objectives for the test. These will be used to focus testing and ensure relevant results. Specifically, the expected security model of the target is discussed, and high impact compromises of the model are identified as objectives.
- **Gather Relevant Documentation:** Tevora works with the Client to acquire IT and business process documentation. Tevora can also take a zero-knowledge approach and attempt to acquire this information during the reconnaissance phase of the test.
- **Determine Level of Access:** Based on the objectives, Tevora and the Client determine if credentials are to be provided by the Client for testing. For the most thorough testing, Tevora will use low-level privileges.
- **Time Estimation:** Tevora determines the estimated time needed to cover the scope for the decided testing types.
- **Role Identification:** Tevora assigns a project lead, technical lead, and assistant technical lead to the test. Tevora's technical will have web services and web application specialists assigned to the project including at least one subject matter expert (SME) on the in-scope technologies.
- **Kickoff Meeting:** Tevora and the Client review the planned scope, discuss the project overview, and propose scheduling.
- **Testing Contact Identification:** Tevora and the Client identify their respective points of contact and determine testing status update intervals. Tevora provides an escalation list to the Client.
- **Incident Handling:** Tevora and the Client agree to a response plan for unexpected issues during testing.
- **Project Checklist:** Tevora ensures that every item for the project is checked prior to beginning the penetration test.

After preparation has been completed, the project checklist reviewed, and scheduling finalized, Tevora will begin the penetration test on the scheduled date.

## Phase 1: Reconnaissance

The first phase of a penetration test is reconnaissance. This phase is conducted to gather information on the target and enumerate potential threat vectors. Tevora performs reconnaissance in a strategic manner that emulates the process of real-world adversaries. This process, called Open Source Intelligence Gathering (**OSINT**), is a multi-level approach that consists of several types of information gathering activities.

**OSINT** is done in three phases: **Passive**, **Semi-Passive**, and **Active**:

- **Passive:** Tevora searches the internet for information that is posted by the Client or their employees. Tevora reviews third-party databases that could contain archived Client or employee information including Google, Shodan, and social networking platforms. Traffic is never sent to the Client during this phase, making the testing difficult to detect.
- **Semi-Passive:** Tevora gathers information on the target using requests disguised as normal internet traffic, including DNS requests, service probes, and analysis of document metadata. Traffic may be sent to the Client but will be difficult to detect.
- **Active:** Tevora uses ping sweeps, port scans, banner grabbing, vulnerability scans, and forced browsing to actively enumerate the Client's attack surface. This is a more aggressive phase of reconnaissance that generates significant amounts of abnormal traffic. Tevora gathers a significant amount of reliable information on the Client's systems during this phase. This phase is most likely to be detected by the Client.

## Phase 2: Threat Mapping

Tevora analyzes the information gathered during the reconnaissance phase to map targets to potential threat vectors. This map is used to enumerate threats to the business and prioritize testing on high-impact targets.

The threat mapping phase closely follows the PTES Standard's threat modeling phase. During threat mapping, Tevora performs the following steps:

- **Gather relevant documentation:** Tevora works with the Client to acquire IT and business process documentation. Tevora can also take a zero-knowledge approach and attempt to acquire this information during the reconnaissance phase.
- **Identify and categorize primary and secondary assets:** Tevora identifies the assets on the in-scope targets and divides them into primary and secondary categories. These are assets that can be reached directly, and assets that can be reached from pivoting, respectively.
- **Identify and categorize threats and threat communities:** Tevora enumerates the potential threats to the in-scope targets and categorizes them by the groups of people (e.g., threat communities) that may execute those threats.
- **Map threat communities against primary and secondary assets:** Tevora maps the categorized threat list to the categorized asset list to determine relevant threats and their potential impact on the business.
- **Cross-reference threat map to test objectives:** Tevora reviews the threat map to identify the impact of potential threats in the context of testing objectives defined during the planning phase.

Tevora uses the output of this phase to enumerate potential threat vectors and prioritize testing on high-impact attack scenarios. This also enables alignment of threat exposure to testing objectives.

## Phase 3: Known Vulnerability Identification

Tevora reviews information gathered during the threat mapping and reconnaissance phases to identify known vulnerabilities. Tevora reviews banners, network and HTTP response signatures, and running services. These are then cross-referenced against vulnerability databases such as Exploit-DB, Rapid7, and CVE.

Tevora takes a multi-assessment approach by analyzing information gathered from both passive and active vulnerability identification:

- **Passive:** Tevora reviews metadata from public documents and archived content in search engines for vulnerability signatures. Additionally, Tevora performs traffic monitoring on the internal network and analyzes network protocols for signatures of vulnerable network services.
- **Active:** Tevora uses vulnerability scanners for automated vulnerability enumeration and augments this with output from port scanners, HTTP responses, SNMP enumeration, NetBIOS enumeration, and more.

After identifying vulnerabilities, Tevora attempts to validate vulnerabilities and prioritize them for exploitation. Tevora researches all discovered vulnerabilities and performs manual testing to check for false positives. Vulnerabilities are cross-referenced against the threat map to identify their impact and potential risk to the business.

## Phase 4: Input Testing

Tevora tests for input validation and injection issues on web application forms. Tevora fuzzes input fields using a combination of manual and automated techniques.

Tests performed include:

- LDAP Injection
- ORM Injection
- Directory Traversal / File Inclusion
- XML Injection
- SSI Injection
- XPath Injection
- IMAP/SMTP Injection
- Code Injection
- OS Commanding
- Buffer Overflow
- Incubated Input Vulnerability Testing
- HTTP Splitting/Smuggling
- SQL /NoSQL Injection

Any discovered input vulnerabilities are categorized, referenced against the threat map, and documented for use in the exploitation phase.

## Phase 5: Application Logic Testing

Tevora uses information gathered during previous phases to analyze the web application's security model. Tevora reviews the application logic at all points in the platform to identify application logic weaknesses which may expose sensitive information or functionality.

Application logic tests make up the bulk of time spent on web application penetration tests. Application logic testing is a primarily manual process, with custom scripts and plugins used to automate testing for certain flaws, such as enumeration. Tevora focuses on high-impact functionality during this phase as well as complex multi-step processes, which are more likely to include dangerous bugs.

Testing performed includes:



The impact of any data leakage, or unauthorized activities discovered during application logic testing are categorized and referenced against the threat map. If relevant, identified issues may be used during the platform exploitation phase.

## Phase 6: Exploitation

During the exploitation phase, Tevora attempts to access the targets enumerated during the threat mapping phase. Tevora reviews discovered vulnerabilities and potentially insecure services to develop an exploitation plan. Tevora then executes this plan in a precision strike against the Client.

Tevora uses publicly available exploits and pursues development of custom and/or “zero-day” exploits for high impact targets or when known vulnerabilities are not discovered.

- **Known Vulnerabilities:** Tevora modifies public exploits to target the Client environment. Public exploits are only acquired from trusted sources such as Exploit-DB and are reviewed before modification and use. Commercial exploitation frameworks are also used during this phase.
- **Unknown Vulnerabilities:** If known vulnerabilities are not found, Tevora takes a zero-day approach. A replica environment is created and Tevora tests the discovered services for previously unknown security issues.
- **Application Layer Vulnerabilities:** If any custom applications are discovered during testing, Tevora will perform application-level assessments as permitted by the timeframe. These tests will be performed according to Tevora’s application testing methodologies.

Tevora delivers payloads during the exploit to gain access to the targets in accordance with testing objectives. Payloads are designed to bypass security measures used by the Client. These will include encoded, packed, encrypted, and custom payloads designed to bypass anti-virus, IPS/IDS systems, and firewalls. These payloads are also used in the post-exploitation phase to pivot the attack to other targets.

## Phase 7: Post-Exploitation

During this phase, Tevora evaluates the impact of the exploitation, tests the Client's internal defenses, and uses the initial exploits to escalate access to additional targets. The following activities are performed during this phase:

- **Establish Persistence:** Tevora establishes secure, persistent access so Tevora may notify the Client of the exploit and the Client can remediate without interrupting post-exploitation activities.
- **Initial Enumeration:** Compromised resources are enumerated for relevant information. User accounts and passwords are extracted for use in pivoting.
- **Pivoting:** Tevora repeats the reconnaissance, threat mapping, vulnerability identification, and exploitation phases on newly accessible targets. Tevora begins the new reconnaissance phase with network analysis and shifts to an internal penetration test methodology. Tevora uses information acquired during previous phases to escalate access to the Client's systems.
- **Target Profiling:** Tevora enumerates data and information on exploited targets.
- **Data Exfiltration:** Based on the purpose of the penetration test, Tevora targets and attempts to extract (or simulate an extraction of) information that is vital to the organization.
- **Cleanup:** When the penetration test is complete, Tevora cleans up all the tools and payloads that were placed in the target's environment.

Post-exploitation is an iterative testing process to continually escalate the attack simulation. Previous steps of the methodology are repeated to assess potential threats from the newly acquired foothold. Additional information about the target may be discovered during this phase such as source code, undocumented endpoints, and additional credentials, which all warrant further testing.

## Phase 8: Reporting

Tevora compiles the findings during the penetration test and organizes them into a final report which is sent to the Client. The report documents each discovered vulnerability, remediation recommendations, and provides an analysis of risk to the business.

Topics covered by the report include:

- Executive Summary
  - People involved
  - Project objective
  - Project scope
- Findings Overview
  - Test results
  - Strategic recommendations
- Technical Summary
  - Scoring of findings
  - Findings summary based on HydraRisk model
  - Detailed summary of each finding
    - CVSS score
    - HydraRisk score
    - Finding description
    - External references
    - Recommended remediation
- Penetration Testing Methodology

The report provides both a detailed technical breakdown and a high-level executive summary, allowing for review by both technical and non-technical staff. The report can be tailored to a Client's needs, including being split into multiple documents. The report is the final deliverable for testing and may go through review and editing phases prior to acceptance. Once the report has been accepted, the project is considered closed unless otherwise stated.



# TEVORA™

Go forward. We've got your back.

Compliance – Enterprise Risk Management – Data Privacy – Security Solutions – Threat Management

