Incident Response & Continuity Plan

This document describes Tekle Holographics' Incident Response (IR) plan and Business Continuity/Disaster Recovery (BC/DR) strategies in detail. It is intended for our customers' security teams and auditors who want to understand how Tekle prepares for and manages incidents that could impact security or availability. Having a well-defined IR plan ensures that if the worst happens, Tekle can respond quickly and effectively, minimizing damage and fulfilling our obligations to customers. Meanwhile, our continuity plan explains how we keep critical services running and recover from large-scale disruptions. This document provides step-by-step insight into these processes, demonstrating Tekle's readiness and resilience.

Incident Response Plan (IRP)

Our Incident Response Plan is a living document that guides our team through the stages of any incident. The plan follows the **NIST 800-61r2** lifecycle: Preparation; Detection & Analysis; Containment, Eradication & Recovery; and Post-Incident Activities. Below we break down how Tekle addresses each stage:

Preparation:

- **Team & Roles:** Tekle has an Incident Response Team (IRT) in place, led by our Security Incident Manager (SIM). The IRT includes members from Security (for analysis and coordination), IT Operations (for system-level issues), Development (for code issues), and Communications (for customer and public communication). We have primary and backup contacts for each role. For example, if the SIM is unavailable, a deputy SIM steps in.

- **Training & Drills:** The IRT undergoes training on IR procedures at least annually. We also conduct incident drills (tabletop exercises and live simulations) twice a year. These drills might simulate, say, a ransomware attack on an internal system or a vulnerability exploitation on our cloud – and we practice our response. Lessons from drills are used to update the IR plan.

- **Incident Classification:** We define severity levels for incidents:
  - *Severity 1 (Critical):* Severe impact or imminent impact on multiple customers or sensitive data (e.g., confirmed breach of customer data, widespread service outage).
  - *Severity 2 (High):* Significant issue but perhaps limited in scope (e.g., breach attempt detected and blocked, or a vulnerability that could lead to breach if not fixed, or a service outage affecting a subset of customers).
  - *Severity 3 (Medium):* Moderate issues like isolated service disruptions, minor security anomalies requiring investigation.
  - *Severity 4 (Low):* Minor incidents, routine malware on an employee workstation, false positives, etc., that don't affect customers materially.

    Each severity has an associated response requirement – for Sev1 and Sev2, we assemble the core IRT within minutes to an hour; for Sev3, within business hours; Sev4 can be handled in normal course.

- **Communication Channels:** We have established secure communication means for the IRT. During an incident, we avoid normal email for sensitive info; instead, we use an encrypted chat channel (e.g., an incident Slack/Teams channel with limited membership and message retention turned off after resolution) or a secure bridge line. We also maintain an on-call contact list (with personal phone numbers as backup) in case incidents happen after hours.

Detection & Analysis:

- **Monitoring Systems:** As noted in other documents, Tekle has multiple monitoring systems (SIEM for log analysis, cloud service monitors, host-based EDR agents, etc.). When these systems detect a potential issue (e.g., an EDR agent flags malware, or monitoring shows unusual traffic, or an application alert triggers), they generate an alert.

- **Incident Ticketing:** All alerts funnel into our incident tracking system (we use a dedicated IR module in our service management tool). An incident ticket is created with initial details. The on-call security analyst reviews the alert and determines if it's a true incident. If confirmed or even suspected with reasonable likelihood, they escalate per the IRP.

- **Incident Declaration:** We have an "Incident Declaration" step – essentially a quick huddle or message where the SIM (or on-call person) declares: "We are now in incident mode for [description] and it's classified as Sev X." This triggers everyone needed to join.

- **Investigation & Analysis:** The IRT analysts begin collecting data:
  - For a security breach, this means preserving relevant logs (we might immediately pull log backups to ensure no loss), capturing system images if a host is compromised, noting all indicators of compromise (IoCs).
  - For a service outage, this means checking system metrics, recent changes (did a deployment just occur?), and pinpointing the fault.
  - We form a hypothesis of what happened and what the impact is. For example, "Web server was serving error pages starting at 10:05 UTC due to a spike in CPU – possible DDoS" or "Multiple customer accounts were locked out, possibly brute-force attempts." We verify what data or services are impacted.
  - Throughout analysis, we document everything in the incident ticket: timelines, actions taken, people involved.

- **Engaging Experts:** If needed, we have external IR partners (security consulting firms) on retainer that we can call for additional forensic expertise or handling a major breach. We also contact our cloud provider's incident response team if their infrastructure is part of the issue (e.g., an Azure outage or suspected platform issue).

- **Customer Impact Assessment:** A crucial part of analysis is determining if customers are affected and how. Are systems down? Was customer data accessed or exfiltrated? We categorize affected customers (specific tenants, or all). This informs our containment and communication.

Containment:

Once we understand the basics of what's happening, we move to contain:

- **Short-Term Containment:** Immediate actions to stop the bleeding. If a server is compromised, isolate it (e.g., remove from network or shut it down if necessary). If an account is compromised, disable or reset it. If malware is spreading, segment network or kill processes.

- **Backup Activation:** For availability incidents, containment might involve failing over to backup systems. For example, if our primary database is malfunctioning, we might activate the standby replica and reroute services to it while we contain the issue on primary.

- **Network Blocks:** We often implement network-level blocks as needed: blocking an IP range that's attacking us, or geofencing if an attack originates from a region where we have no customers.

- **Communication during Containment:** We keep the team and stakeholders updated. For a high-severity incident, we may send an initial note to affected customers like "We are aware of an issue causing [symptom]. Containment actions are underway. More info to follow." This lets them know we're actively responding even if we don't have full details yet.

Eradication:

After immediate containment, we ensure the threat is eradicated:

- If it's malware: ensure all infected systems are cleaned (rebuild machines, if necessary, that's often safest).

- If it's a compromised user account: make sure we identify how they got in (phish? guessing? and address that, e.g., tighten MFA or fix a vulnerability they exploited) and check they didn't create any backdoor accounts or leave any malicious jobs/scripts. We'll scour logs for any additional IoCs.

- If it's vulnerability exploitation: once contained (like turned off that feature or isolated that service), we apply the fix or patch to eliminate the vuln. This could involve deploying a code hotfix or configuration change.

- We confirm through testing that the root cause is gone. For example, after cleaning a webshell from a server, we run scans to ensure no other webshells or malicious files exist on other servers.

Recovery:

Now we bring services back to normal operation carefully:

- Restore systems from clean backups if needed (if a database was corrupted or data encrypted by ransomware and we have to restore from backup, for instance).
- Ensure systems are patched and hardened before reconnecting them.
- Monitor them closely as they go live to ensure no reoccurrence.
- For outages, we communicate to customers when service is restored and any precautions (like "we rolled back to a previous version" or "you may need to restart your client application" etc.).
- If data was lost or altered, recovery includes data reconciliation. For example, if some transactions didn't get processed during an outage, we might re-run queued tasks to catch up.

Notification & Communication:

Parallel to the above, communication with stakeholders (especially customers) is handled:

- **Internal Communication:** Our execs and relevant staff are kept updated through an incident channel. If it's serious, the CEO or CTO may get involved in decisions (particularly around public communication or major expenditures like offering credit monitoring if personal data breach).
- **Customer Communication:** We have templates for incident notification to customers. Our policy is transparency without causing unnecessary alarm. A typical breach notification will include what happened, what data or services are impacted, what we have done about it, and what we recommend the customer do (e.g., reset passwords if credentials were compromised, or simply be aware of the downtime).
- We respect any regulatory timelines (GDPR's 72-hour breach notification requirement to authorities and controllers, for instance). We coordinate with customers on notifications – often we, as a processor, prepare breach info for the customer (controller) to notify users or authorities, unless we are required to notify directly.

- We might set up a dedicated email hotline or conference call for affected customers to ask questions.

- **Public Communication:** If the incident is likely to become public (widespread outage or known breach), we prepare public statements (often via our website status page or press if needed). We ensure not to disclose sensitive details that could aid attackers, but enough to maintain trust.

Post-Incident Activities:

Once the incident is resolved and things are stable:

- **Post-Mortem Meeting:** Within a few days (once emergency stress has passed), the IRT and management hold a post-mortem. We review the timeline, identify root causes (both technical and process), evaluate what went well or poorly.

- **Action Items:** We derive specific actions: e.g., "Implement additional monitoring for X," "Update password policy," "Train staff on social engineering based on this phish," "Improve our backup testing frequency," or "Change a vendor if they were at fault." We assign owners and deadlines to these actions.

- **Report:** We compile an Incident Report document that records all facts: what happened, how we responded, what data was affected, and what we're doing to prevent it. This report (sanitized if necessary) can be shared with customers, especially for significant breaches, to demonstrate our accountability and improvements.

- **Lessons to Training:** If the incident revealed a knowledge gap, we feed that into our next training. If it showed a deficiency in the IR plan, we update the plan. For example, if communications were confused, we might adjust our severity definitions or who must be paged for certain incidents.

Our IRP is designed to be thorough but also flexible – no two incidents are exactly alike, so our team is trained to adapt while still following the general structure. The ultimate goal is to protect our customers' data and our service integrity above all, even if it means taking aggressive actions (like isolating systems or bringing down parts of service briefly to stop a breach). We prioritize security over short-term convenience in those moments.

Incident Scenarios and Responses

To illustrate, here are brief examples of how we'd handle specific incident scenarios:

- **Scenario 1: Ransomware Detected on an Employee Computer** – Our EDR flags encryption activity on a corporate laptop. Containment: laptop is cut from network by IT, user informed not to turn it off (to preserve evidence). IR identifies if any spread – check file servers, etc. It seems contained to that one device. Eradication: reimage laptop, restore user's files from backups (we keep user data in cloud storage which has versioning). Since it was caught early, no impact on production or customers. We investigate root cause (user clicked a phishing link). Post-incident: do a phishing awareness refresher for staff, maybe implement stricter email filtering.

- **Scenario 2: Customer Data Breach** – Unusual database queries are observed on our cloud database – looks like someone is extracting data in bulk. Analysis shows a compromised admin API token being used. Containment: Immediately revoke that token (and all tokens for that account), temporarily firewall the database from that API to stop the flow. Eradication: Identify how they got the token – turns out an admin's credentials were phished; we reset all admin creds and improve 2FA enforcement. Assess what data was taken – if logs show they pulled records from, say, one particular table of user info, we identify those records (which customers impacted). Recovery: no system down, but we might refresh any exposed secrets (if any API keys in data, etc.). Notification: Inform affected customers that certain data (e.g., usernames and emails, maybe design model names, etc. – whatever it was) was accessed by an unauthorized party, with apologies and actions taken. Post-mortem: tighten access controls, maybe reduce the data accessible with a single token (paging design review perhaps to limit how much data an admin can pull in one go, or implement anomaly detection to catch large exports sooner).

- **Scenario 3: Major Cloud Outage** – Azure experiences a regional outage taking down our production environment. Our continuity plan kicks in: Containment is mostly waiting for Azure, but we decide to failover to a secondary region which is still up. We spin up our DR environment (which we have on warm standby), restore latest data backups if cross-region replication was incomplete up to the minute. We notify customers of an outage and that we are performing DR failover. Recovery: in a couple hours service is running in secondary region. We later fail back or adjust architecture to be multi-region. Post-incident: review if failover was as smooth as expected, update any runbooks, consider keeping multi-region active-active to avoid downtime.

- **Scenario 4: Vulnerability in Tekle Software (0-day)** – A security researcher contacts us about a vulnerability they found in our API that could allow unauthorized data access under certain conditions. This triggers our IR process (even though it's not an active incident, it's a potential one). Containment: we verify it and deploy an immediate hotfix patch to cloud (and prepare patch for on-prem customers). If the vuln is already being exploited, we'd check logs and possibly disable the affected feature until patched. Eradication: patch is applied, vulnerability closed. Communication: Possibly issue a security advisory to all customers (especially if on-prem deployments need patching) describing the issue in general and urging upgrade to patched version. Work with the researcher for responsible disclosure (maybe give them credit and/or bug bounty). Post-incident: improve our testing to catch similar vuln, thank researcher publicly.

These scenarios show how having a plan enables quick, effective action tailored to each situation.

Business Continuity and Disaster Recovery (BC/DR)

Tekle's Business Continuity plan ensures that critical business functions can continue or quickly resume in the event of major disruptions, while Disaster Recovery focuses on restoring IT systems after catastrophic failures.

Business Impact Analysis (BIA): We have identified our critical services and processes. The most critical is our **production cloud service** (holographic collaboration platform) because downtime directly affects customers. Others include: our R&D (so we can produce fixes), our support operations (to assist customers during issues), and internal functions like finance (less immediate but included in continuity for long-term outages). For each, we have set Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO):

- Production Cloud Service: RTO = 4 hours (goal to recover service within that time in DR scenario), RPO = 1 hour (goal to not lose more than 1 hour of data transactions).

- Internal Development: RTO = 24-48 hours (since if our dev tools go down, not immediately impacting customers, but we need to recover to continue work), RPO = ideally near 0 for code (we use distributed source control so code is on dev machines as well as cloud).

- Support systems (ticketing, comms): RTO = 4 hours, RPO = 0 (we have cloud-based support tools with their own SLAs, and our backup is direct email/phone if the tool is down).

Continuity Strategies:

- For production, we maintain redundancy as mentioned: multiple instances, failover capability to secondary region or cloud. We have backup infrastructure contracts if primary cloud is completely unreachable (e.g., if Azure had a massive outage, we could port to AWS in worst case, though that would be slower – that's beyond RTO usually, so realistically if both primary and secondary in Azure fail, it's extreme).

- Our DR plan for cloud is documented step-by-step (how to restore databases from backup, how to repoint DNS, etc.). We test parts of this plan quarterly (for instance, test restoring a backup to a new environment and connecting a test client).

- Key data is backed up off-site: in addition to cloud replication, we take encrypted snapshots of databases and store them in a separate cloud storage (and a copy off the cloud provider – e.g., Azure backups also copied to an AWS bucket or our on-prem NAS, to guard against cloud account issues).

- Our office or people continuity: Tekle's workforce has shifted to be capable of remote work. If our main office is inaccessible (fire, etc.), employees can work from home securely (we use VPNs to access internal resources). We ensure critical staff have laptops with necessary tools. Our phone system is cloud-based so it's reachable anywhere.

Disaster Recovery Plan Implementation:

- We keep an updated "Runbook" for DR that authorized personnel can execute. It includes credentials (securely stored) for DR infrastructure, and instructions like:

    1. Declare Disaster (criteria: primary site down > X hours or likely to be down > Y hours).

    2. Convene DR team (which overlaps IR team but also IT infra folks).

3. If using secondary region: promote secondary DB to primary, update config, redirect traffic via DNS or load balancer to secondary.

4. If secondary was not already running: spin up environment from infrastructure as code templates, restore data from backup, etc.
5. Test basic functionality with a small user group, then announce to all clear that service is up.

- We emphasize making decisions based on estimated outage length. If primary might recover quickly, sometimes better to wait; but if uncertain, we err on activating DR to minimize customer downtime.

Communication in BC events: We have a crisis communication plan. For example, if something like a natural disaster affects our region (maybe making support slower), we notify customers that "Due to [event], response times may be longer, but our cloud systems are unaffected" etc. If systems are affected, we use status pages and direct email to share updates regularly (at least every hour or two for a significant outage, even if just "still working on it").

Alternate Site and Resources: Tekle maintains accounts with alternate cloud providers in case of extreme need. We also keep some critical documentation (like the IR/DR plans, contact lists) in hard copy or on devices of key personnel so they can access instructions if digital systems are down. In our worst-case scenario planning, if both our primary and backup fail, we would communicate openly with customers, perhaps providing a degraded service or workarounds (for example, if our cloud was completely down, clients with on-prem capability might switch temporarily to offline mode to continue using their devices standalone until we restore central services).

Restoration and Return to Normal: After a disaster scenario, once the primary environment is back, we plan a controlled return (either stay in new environment if stable, or sync data back and switch). We perform a full review post-DR activation too, looking to improve our DR switchover times or processes.

Logging, Auditing, and Continuous Improvement

We covered how logs are used in IR. Here we add how logs and audits help ensure continuity:

- Tekle conducts quarterly **restore tests** (we pick a backup and simulate a restore to ensure backups are viable).
- We audit user access to backups (only a couple of ops engineers can access backup storage).
- We log any use of emergency accounts or procedures (like if someone uses the break-glass admin account in a disaster, it's logged and reviewed).
- Every year, we audit the entire IR/BCDR process via either an internal audit or external consultant to ensure nothing is outdated (for instance, as infrastructure evolves, are plans updated? Are contact lists current?).

Tekle Holographics is prepared to face security incidents and operational disruptions with a structured and practiced plan. Our Incident Response process ensures we can react swiftly to contain and resolve threats, keeping stakeholders informed along the way. Our Business Continuity and Disaster Recovery strategies provide confidence that even in extreme scenarios, we can restore key services and support our customers through the event. We continuously refine these plans through drills and real-world lessons because trust isn't just about preventing incidents – it's also about handling them professionally and transparently when they occur.

We share this plan with you not only to be transparent, but also to encourage coordination. In the event of a broader incident that affects both Tekle and your organization, we will work together with your incident management team for a unified response. Our commitment is that you will never be left in the dark; we will be a reliable partner in any crisis. By understanding our procedures, you can better integrate us into your own incident and continuity planning.

For any questions about our incident response or continuity capabilities, or if you would like to perform joint incident exercises, please reach out to our security team. Collaboration and preparation are the keys to resiliency, and we value the opportunity to strengthen both with our customers.