

Jira Continuity Card

A starting point for your own binder. Page 1 & 2 is the card you fill in. Page 3 is the reference. The pages after that are worked examples, one per failure mode, so you can see what a completed card looks like before you write yours.

These should absolutely be part of your preparation list, though to be clear this list is not exhaustive, and your details will (not may) vary from what I have.

The Card

Defining our yellow: what we can still do while the recovery is happening. Print it. Keep a copy in two buildings.

Organization	Site / Instance	Doc owner	Last reviewed

Fields marked [NOT THE ADMIN] need a person outside IT to fill in or sign off. If those are blank, this plan will not work on the day you need it.

What we set up in advance

The thing that turns a red event into a yellow one. It has to exist before the outage.

Who can declare degraded mode [NOT THE ADMIN]

Name a role, not a person. This is a business decision, not a technical one.

--

How we announce it [NOT THE ADMIN]

Must not require the system that is down.

--

Check order

The sequence we walk to establish whose outage this actually is, before we act.

1	
2	
3	
4	

What keeps running: the yellow menu [NOT THE ADMIN]

Capability	Fallback while degraded	Path back into Jira

Who we call [NOT THE ADMIN]

Failure type	Team or vendor	Name and number	Escalation

Seven Ways You Lose Jira

Not every outage is the same outage. Know which one you are in before you act.

#	Failure mode	How you detect it	Who actually fixes it
1	Nobody can authenticate	IdP outage, expired SAML certificate, an auth policy change, SCIM deprovisioning error.	Identity or IAM team, or whoever owns the IdP.
2	Nobody can reach it	ISP or network failure, proxy misconfiguration, DNS, an IP allowlist change.	Network or IT infrastructure team.
3	Work cannot get in	Service desk email intake stops accepting or forwarding requests. Nothing looks broken.	Mail flow admin and the service desk owner.
4	A capability quietly disappears	Marketplace vendor outage, a lapsed app license, an app update that changes behavior.	The app vendor, and whoever owns that procurement relationship.
5	The instance is healthy but wrong	A bad permission scheme, a deleted field, a broken workflow transition, a runaway automation rule.	The Jira admin. This is the one case that is genuinely yours.
6	You lose the keys	Billing lapse, or the last org admin leaves without handing off access.	Finance, procurement and org leadership. Not IT.
7	Atlassian is actually down	Status page confirms a platform incident. Note the page can lag your users.	Atlassian. Nothing for you to fix, which is exactly when page one earns its keep.

Of these seven, exactly one is Atlassian's to fix. The other six are yours, and most of them are not the admin's alone.

1. Nobody can authenticate

Atlassian's status page is a wall of green. Every service healthy, your data intact. And not one person can log in.

What we set up in advance. A break-glass org admin account that does not authenticate through the IdP, with credentials in the password safe rather than in Confluence, because Confluence is behind the login that just broke.

Check order

1. Sign in with the break-glass account. If that works, Atlassian is fine.
2. Check the Atlassian status page to rule out a platform incident.
3. Check the IdP status page and admin console.
4. Establish scope: everyone, or only one authentication policy?

How we announce it

SMS tree, then the standing phone bridge. Not email and not Confluence, because you cannot announce an authentication outage using a tool that requires authentication.

What keeps running

Capability	Fallback while degraded	Path back into Jira
Incident intake	On-call phone number and the shared mailbox	Raised as work items in priority order once green, timestamps preserved
Emergency change approval	Named approver by phone, decision written down at the time	Change record created retroactively with the approver named

Who we call

Failure type	Team or vendor	Name and number	Escalation
Authentication	Identity / IAM		IdP vendor support, then their severity one line

2. Nobody can reach it

Jira is up. Atlassian is fine. Your office cannot get to it.

What we set up in advance. A phone on cellular as the standard first test, and a current record of who owns the IP allowlist and when it last changed.

Check order

1. Load the site on a phone with wifi off. Thirty seconds, and you know whether it is us or the internet.
2. If cellular works, the problem is inside our network. If it does not, check the Atlassian status page.
3. Ask when the allowlist, proxy or firewall last changed, and who changed it.
4. Test from a second office or over VPN to establish blast radius.

How we announce it

SMS or phone. Email and chat may share the pipe that just died, so the announcement has to reach handsets directly.

What keeps running

Capability	Fallback while degraded	Path back into Jira
Critical work	Approved cellular or personal devices, if policy allows. That decision is made in advance, not at 9 a.m.	Normal, once the network path returns
Incident intake	On-call phone number	Raised as work items once reachable

Who we call

Failure type	Team or vendor	Name and number	Escalation
Network path	Network / infrastructure		ISP account team, circuit ID to hand

3. Work cannot get in

Nothing looks broken. Intake stopped, Jira is healthy, and customers are sending requests into a void while believing they filed them.

What we set up in advance. An alert that fires when intake volume hits zero during business hours, plus a test address we mail from outside the company. This failure is silent, so the preparation is detection rather than response.

Check order

1. Send a test message from an external address and watch for it.
2. Establish whether it reached the mailbox but was not processed, or never arrived at all.
3. Check the mail handler, forwarding rules and any recent mail platform change.
4. Check Atlassian status for the mail handling component specifically.

How we announce it

Customers first, staff second. This is the one failure where the affected people are outside the organization, so an internal phone tree does nothing. Portal banner plus whatever channel customers actually watch.

What keeps running

Capability	Fallback while degraded	Path back into Jira
Customer request intake	Published alternate address and phone number, logged manually as they arrive	Entered as work items with original received times
The gap	Named owner asks senders what they submitted during the outage window	Reconciled list raised and acknowledged to each customer

Who we call

Failure type	Team or vendor	Name and number	Escalation
Mail flow	Messaging admin		Service desk owner, then Atlassian support

4. A capability quietly disappears

The instance is fine. One workflow is not.

What we set up in advance. A written map of which app supports which business capability, renewal dates on a team calendar rather than a person's, and a note of which apps update automatically.

Check order

1. Is the app itself erroring, or the workflow that depends on it?
2. Check the vendor's status page, which is not Atlassian's.
3. Check the license state in the admin console.
4. Check whether the app updated recently. A vendor's release is your incident.

How we announce it

Narrow. Only the teams whose workflow actually stopped. A company-wide notice for one app makes the next notice easier to ignore.

What keeps running

Capability	Fallback while degraded	Path back into Jira
The affected workflow	Manual version, if one exists. Sometimes the honest answer is that it does not, and that is worth knowing before today.	Backfilled once the app returns

Who we call

Failure type	Team or vendor	Name and number	Escalation
App or vendor	Vendor support		Procurement owner. Know your support tier: it means something different at a four person vendor.

5. The instance is healthy but wrong

Every service is green. Your data is intact. And something you did on Thursday is now doing damage at scale.

What we set up in advance. Every config change recorded with date, author and scope, somewhere that is not your memory. A sandbox habit for anything touching schemes or automation. This is the one mode where you own the cause and the fix, and the product still gives you no undo.

Check order

1. What changed in the last 24 hours, and who changed it?
2. Audit log for the affected object, assuming your plan records those events.
3. Blast radius: how many work items has it touched, and can you list them?
4. Can this be reverted, or only corrected forward?

How we announce it

You say it, yourself, immediately, with scope. The instinct is to fix it quietly before anyone notices. The admins who announce their own mistakes early are the ones who still have credibility a year later.

What keeps running

Capability	Fallback while degraded	Path back into Jira
The broken workflow	Disable the offending rule, freeze further config change, manual approval path	Corrected items reprocessed in a controlled batch

Who we call

Failure type	Team or vendor	Name and number	Escalation
Self inflicted config	You		Atlassian support, and only for a site level restore

6. You lose the keys

This one usually arrives with a warning nobody read.

What we set up in advance. At least two active org admins, both of whom have logged in this quarter. A billing contact and payment method owned by a role rather than a person. Renewal and true-up dates on a team calendar. An offboarding step that verifies remaining org admin access before an account is disabled, not after.

Check order

1. Is this a billing problem or an access problem?
2. What is the deadline before functionality actually degrades?
3. Who still holds org admin, and can they log in right now?
4. Who owns the commercial relationship with Atlassian?

How we announce it

Leadership and procurement, not users, until the scope and the deadline are known. This is a governance failure that becomes an outage, and announcing it as an outage first invites the wrong response.

What keeps running

Capability	Fallback while degraded	Path back into Jira
Everything, briefly	This is a countdown rather than an outage. Use the runway to fix it rather than to plan around it.	Not applicable if resolved before the deadline

Who we call

Failure type	Team or vendor	Name and number	Escalation
Billing or ownership	Finance / procurement		Atlassian billing support, then your reseller or partner

7. Atlassian is actually down

No lever, no fix, no escalation that changes the outcome. Whatever you built for the other six is all you have, which is why this is the mode that tests whether the card is real.

What we set up in advance. A subscription to Atlassian status notifications, from before today, and a limited menu that is already written. There is no version of this morning where you invent one.

Check order

1. Atlassian status page.
2. Independent confirmation from users or a third party monitor, because the status page can lag what your people are experiencing.
3. Scope: which products, which regions, which of your sites.
4. Their stated ETA, if there is one, and when they will next update.

How we announce it

Immediately, and with the menu. Not "we are investigating," which gives nobody anything to act on. Say what still works, how to do it, and who to ask.

What keeps running

Capability	Fallback while degraded	Path back into Jira
The full yellow menu	Everything on the card. This is the event it was written for.	Reconciliation run by the named owner once service returns
Everything else	Waits	Resumes

Who we call

Failure type	Team or vendor	Name and number	Escalation
Platform incident	Nobody who can fix it		Atlassian enterprise incident contact, if you have one