FREE ISO 20000 AUDIT TOOLKIT

ISO 20000 Gap Analysis Checklist

30 audit-ready checks to find service management gaps before the next outage does

Assess supplier, continuity, capacity, incident and service level management against real cloud outage scenarios.

How to Use This Checklist

A major cloud outage is one of the most honest tests of an IT service management system. Policies, plans and contracts that look complete during a document review are suddenly put under real pressure: Did the supplier escalate? Did the backup environment work? Did customers hear from you before they heard from social media?

This checklist turns those outage lessons into **30 audit-ready checks** across the five ISO/IEC 20000-1:2018 processes that cloud disruptions expose most often. For each check you will find what to look for, the evidence to request, the red flags that signal a gap and a worked example from a realistic outage.

Who it is for

- **ISO 20000 internal auditors and lead auditors** planning or conducting audits of cloud-dependent services.
- **ITSM and service delivery managers** who want to know where their service management system would fail under pressure.

- **Compliance and risk teams** preparing evidence for certification or surveillance audits.
- **Organisations preparing for ISO 20000 certification** that want a practical gap analysis before the certification body arrives.

When to use it

- **Before a certification or surveillance audit**, to close gaps while there is still time.
- **After a major outage or near miss**, while the evidence and lessons are fresh.
- **During the annual management review** of the service management system (SMS).
- **When onboarding a critical cloud supplier**, to make sure the arrangements are right from day one.

The four steps

1. **Choose your scope.** Assess one critical service (for example, the customer portal) or the whole SMS. A focused scope produces

sharper findings; start with your most business-critical, cloud-dependent service.

2. **Rate each check.** Use the rating scale below. Rate what you can prove, not what you believe to be true.

3. **Record the evidence.** Note the document, record, report or interview that supports each rating. This becomes your audit trail.

4. **Summarise and act.** Total your ratings on the Gap Summary page, calculate readiness and build a prioritised action plan.

Rating scale

Rating	Meaning	Example
C Conforming	Defined, implemented and supported by evidence that it works	Continuity plan exists, was tested last quarter and test results are recorded
P Partial	Documented but not fully implemented, or implemented without evidence	Continuity plan exists but has never been tested

G Gap	Missing, or clearly not working	No continuity plan for the loss of the primary cloud region
N/A Not applicable	Outside the scope of this assessment	On-premises-only check when the service is fully cloud-hosted

WATCH OUT | The "documented vs implemented" rule

A check is **Conforming** only when there is evidence that the process actually works. A well-written policy with no records, tests or reviews behind it is **Partial** at best.

This is the single most important distinction in an ISO 20000 audit, and it is exactly what major outages expose: organisations rarely fail because a document is missing; they fail because the document was never put into practice.

EXAMPLE | The outage scenario used throughout this checklist

To show how each check works in practice, every section refers to the same realistic scenario:

- **Organisation:** Acme Billing, a mid-size SaaS company providing invoicing and payment software to 2,000 business customers.

- **Set-up:** its customer portal and payment API are hosted in a **single region** of a major public cloud provider.
- **Incident:** a regional outage at the provider lasts **4 hours** on a weekday morning. Customers cannot log in, raise invoices or process payments.
- **Aftermath:** support receives over 600 tickets, two large customers escalate to account managers and the incident report records only "Cloud provider outage, 4h 07m, resolved by provider".

Process 1: Supplier Management

ISO/IEC 20000-1:2018 Clause 8.3.4 • Checks 1–6

Almost every modern IT service depends on external suppliers: cloud infrastructure, SaaS platforms, CDNs, payment gateways and managed service providers. ISO 20000 requires the organisation to manage these suppliers so that the services it delivers to its own customers meet their requirements.

The organisation may not control the supplier's technology, but it is fully responsible for how it selects, contracts, monitors and plans around that supplier.

> **OUTAGE LENS | Your supplier's outage becomes your service failure**
>
> When the cloud provider went down, Acme Billing's customers did not blame the provider; they blamed Acme. Supplier management is where an auditor checks whether the organisation understood and managed that dependency before the outage, not after.

1. Critical cloud suppliers are identified and risk-assessed

The organisation should know which suppliers its critical services depend on, and what would happen to each service if that supplier failed. Without this map, the impact of an outage is discovered during the outage.

Evidence to request

- A supplier register that marks critical suppliers and links them to the services they support.
- Supplier risk assessments covering availability, concentration and exit risk.

- Evidence that the register and assessments are reviewed (for example, annually or after major changes).

Red flags

- Supplier list exists but does not link suppliers to services.
- Cloud infrastructure provider is missing from the risk register because "it's too big to fail".

> **EXAMPLE | Acme Billing**
>
> Acme's supplier register listed the cloud provider as "Hosting – standard" with no link to the portal or payment API and no risk assessment. **Rating: G.**

2. Supplier evaluation covered resilience before onboarding

Before a critical supplier is selected, the organisation should evaluate more than price and features. Disaster recovery arrangements, incident communication, escalation routes and multi-region options all determine how an outage will play out.

Evidence to request

- Supplier selection or due-diligence records for critical suppliers.

- Evaluation criteria that include resilience, DR, incident communication and support model.
- Review of the supplier's published resilience documentation or independent assurance reports.

Red flags

- Selection was based only on cost or a technical proof of concept.
- No one can explain how the supplier communicates during an incident.

> **EXAMPLE | Acme Billing**
>
> The cloud provider was chosen five years ago by a former CTO. A cost comparison exists, but there is no record that DR or incident communication was evaluated. **Rating: P** (some evaluation, but not of resilience).

3. Contracts define availability, support, escalation and notification

The contract or agreement with each critical supplier should state what the organisation can expect during a disruption: availability commitments,

support response times, escalation contacts and how quickly the supplier will notify the customer of an incident.

Evidence to request

- Signed contracts or terms of service for critical suppliers.
- Documented support tier, escalation contacts and incident notification commitments.
- Mapping of supplier obligations to the organisation's own service requirements.

Red flags

- The organisation is on a basic support tier with no escalation route for a critical service.
- Nobody knows where the supplier's status page or incident notifications are published.

> **EXAMPLE | Acme Billing**
>
> Acme used the provider's standard online terms on a basic support plan. During the outage, the operations team relied on social media posts to understand what was happening. **Rating: G.**

4. Supplier performance is monitored and reviewed

Agreed targets are only useful if someone checks them. The organisation should monitor supplier performance against commitments and review it regularly with the supplier or internally.

Evidence to request

- Supplier performance reports or dashboards (availability, incidents, response times).
- Minutes or records of supplier review meetings.
- Actions raised from supplier reviews and evidence they were followed up.

Red flags

- Supplier performance is only discussed when something goes wrong.
- The organisation relies on the supplier's own reporting without any independent monitoring.

> **EXAMPLE | Acme Billing**

> Acme's monitoring tool tracks the provider's uptime, and a quarterly vendor review is held, but reviews have no agenda or recorded actions. **Rating: P.**

5. Single-supplier and single-region dependency is documented and mitigated

Concentration risk is one of the main lessons of major cloud outages. If a critical service depends on one provider in one region, the organisation should have consciously accepted that risk or taken steps to reduce it.

Evidence to request

- Architecture documentation showing where critical services run.
- Risk register entry for supplier or region concentration, with an owner and a treatment decision.
- Evidence of mitigation: multi-region deployment, standby environment or documented risk acceptance by management.

Red flags

- Single-region dependency was never discussed with management.
- A "multi-region" design exists on paper but has never been built.

> **EXAMPLE | Acme Billing**
>
> The single-region set-up was never formally assessed or accepted by management. **Rating: G.**

6. Supplier-caused incidents feed back into supplier reviews

Each supplier-related incident is evidence about that supplier's performance. The organisation should use these incidents to update supplier risk, renegotiate terms or change arrangements.

Evidence to request

- Incident records tagged with the supplier involved.
- Supplier review records that discuss recent incidents.
- Changes made after supplier incidents (contract updates, new support tier, architecture changes).

Red flags

- Supplier incidents are closed as "external cause, no action".
- The same supplier causes repeated incidents with no change in arrangements.

EXAMPLE | Acme Billing

After the outage, Acme opened a ticket to upgrade the support plan and added the incident to the next vendor review agenda. **Rating: C** (for the post-outage response; check again at the next audit).

GOOD EVIDENCE | What good looks like for Supplier Management

- A supplier register that links each critical supplier to the services it supports, with a current risk rating.
- Contracts that specify availability, support response, escalation and incident notification.
- Regular, minuted supplier reviews with actions that are tracked to closure.
- A documented management decision on concentration risk.

Process 1 rating table

#	Check	Rating	Evidence reviewed	Notes

1	Critical cloud suppliers identified and risk-assessed	☐C ☐P ☐G ☐N/A		
2	Supplier evaluation covered resilience before onboarding	☐C ☐P ☐G ☐N/A		
3	Contracts define availability, support, escalation and notification	☐C ☐P ☐G ☐N/A		
4	Supplier performance monitored and reviewed	☐C ☐P ☐G ☐N/A		
5	Single-supplier / single-region dependency documented and mitigated	☐C ☐P ☐G ☐N/A		
6	Supplier-caused incidents feed back into supplier reviews	☐C ☐P ☐G ☐N/A		

Process 2: IT Service Continuity Management

ISO/IEC 20000-1:2018 Clause 8.7.2 • Checks 7–12

IT service continuity management prepares the organisation to keep critical services running, or restore them within agreed timeframes, when a major disruption happens. ISO 20000 expects continuity requirements to be identified, plans to be in place and those plans to be tested.

Continuity is the area where the gap between "documented" and "implemented" is usually widest.

> **OUTAGE LENS | An outage is a live continuity test**
>
> A real outage tests the continuity plan without warning. Auditors should ask what the plan said would happen, what actually happened, and why there was a difference.

7. Critical services identified and business impact assessed

Continuity starts with knowing which services matter most and what a disruption would cost the business. A business impact assessment (BIA) records the operational, financial, customer and regulatory impact of losing each service over time.

Evidence to request

- A list of critical services approved by management.
- Business impact assessments for critical services, with impact over time (1 hour, 4 hours, 1 day).
- Evidence that the BIA is reviewed when services or business needs change.

Red flags

- Every service is labelled "critical", so nothing is prioritised.
- The BIA was written for a previous version of the service architecture.

> **EXAMPLE | Acme Billing**
>
> Acme had a BIA for the customer portal, but it was written before the payment API was launched and did not include it. **Rating: P.**

8. Recovery objectives (RTO/RPO) defined for each critical service

The recovery time objective (RTO) is the maximum acceptable time to restore a service; the recovery point objective (RPO) is the maximum acceptable data loss. Without them, nobody can say whether a 4-hour outage was acceptable.

Evidence to request

- Documented RTO and RPO for each critical service, agreed with the business.
- Evidence that technical recovery arrangements can actually meet those objectives.

Red flags

- RTO and RPO were set by IT alone without business agreement.
- The RTO is shorter than the time the recovery procedure actually takes.

> **EXAMPLE | Acme Billing**
>
> The portal's RTO was documented as 1 hour. The actual outage lasted 4 hours with no way to recover faster, because recovery depended entirely on the provider. **Rating: P** (defined but not achievable).

9. Continuity plans cover loss of a cloud provider or region

Many continuity plans were written for on-premises scenarios such as data centre fire or power loss. Plans for cloud-dependent services must

specifically address what happens if a provider, region or key managed service becomes unavailable.

Evidence to request

- Continuity plans that include provider or region failure scenarios.
- Defined decision criteria for when to invoke the plan (for example, "provider outage > 30 minutes").
- Named roles and contact lists for plan invocation.

Red flags

- The plan assumes the cloud provider will always recover quickly.
- No one knows who has the authority to invoke the plan.

> **EXAMPLE | Acme Billing**
>
> Acme's continuity plan covered office loss and staff unavailability but contained no cloud region failure scenario. **Rating: G.**

10. Backup or alternative environments have enough capacity and dependencies

A standby environment is only useful if it can actually carry the load. Many organisations discover during an outage that their backup environment is under-sized or missing a dependency such as DNS, identity, secrets or a third-party API.

Evidence to request

- Architecture documentation for the standby or alternative environment.
- Capacity sizing for the standby environment compared with production.
- Dependency mapping showing that all required components are available in the alternative environment.

Red flags

- The standby environment has not been updated since the last major release.
- Identity, DNS or secrets management live only in the failed region.

> **EXAMPLE | Acme Billing**

> Acme had no standby environment for the portal. **Rating: G.** (An organisation with a standby that had never been load-tested would typically be rated **P**.)

11. Continuity plans tested under realistic conditions

ISO 20000 expects continuity plans to be tested. A desk review of the document is a start, but a realistic test, such as a failover exercise or a simulated region outage, provides the evidence that the plan works.

Evidence to request

- A continuity test schedule.
- Test reports with scenario, participants, results and issues found.
- Evidence that issues found in tests were corrected.

Red flags

- The last test was a tabletop discussion more than two years ago.
- Tests are always run in a simplified environment that does not resemble production.

> **EXAMPLE | Acme Billing**
>
> A tabletop exercise was held 14 months ago, covering office loss only. **Rating: P.**

12. Lessons from real disruptions update the plans

Every real disruption is a free continuity test. The organisation should review what happened against the plan and update the plan, BIA and recovery arrangements accordingly.

Evidence to request

- Post-incident reviews that compare actual response with the continuity plan.
- Version history showing the plan was updated after major incidents.
- Actions from reviews tracked to closure.

Red flags

- The continuity plan's version date is older than the last major outage.
- Post-incident actions are recorded but never reach the continuity plan owner.

EXAMPLE | Acme Billing

At the time of the audit, the plan had not yet been updated after the outage, but an action was logged with a due date. **Rating: P.**

GOOD EVIDENCE | What good looks like for Service Continuity

- Approved list of critical services with current BIAs and business-agreed RTO/RPO.
- Continuity plans that explicitly cover provider and region failure.
- Realistic tests at least annually, with results and corrective actions recorded.
- Plans updated after every major incident.

Process 2 rating table

#	Check	Rating	Evidence reviewed	Notes
7	Critical services identified and business impact assessed	☐C ☐P ☐G ☐N/A		

8	Recovery objectives (RTO/RPO) defined	☐C ☐P ☐G ☐N/A		
9	Plans cover loss of a cloud provider or region	☐C ☐P ☐G ☐N/A		
10	Alternative environments have enough capacity and dependencies	☐C ☐P ☐G ☐N/A		
11	Plans tested under realistic conditions	☐C ☐P ☐G ☐N/A		
12	Lessons from real disruptions update the plans	☐C ☐P ☐G ☐N/A		

Process 3: Capacity Management

ISO/IEC 20000-1:2018 Clause 8.4.3 • Checks 13–18

Capacity management ensures the organisation has enough resources, including compute, storage, network, licences and people, to meet current

and future service demand. In the cloud, capacity can scale quickly, but only within the limits, quotas and configurations the organisation has set.

> **OUTAGE LENS | Not every outage starts at the provider**
>
> Some of the most damaging "outages" are self-inflicted: a promotion that triples traffic, an auto-scaling limit nobody remembered, or a database that runs out of connections at peak. Capacity management is where auditors check whether demand was anticipated.

13. Capacity requirements and forecasts exist

The organisation should understand current resource needs and forecast future demand based on business plans, seasonal patterns and growth.

Evidence to request

- A capacity plan covering current usage and forecasts for critical services.
- Input from business plans (marketing campaigns, new customers, product launches).
- Evidence the plan is reviewed at least annually.

Red flags

- Forecasting is based only on last year's usage, ignoring business plans.
- Marketing campaigns are not shared with the IT team in advance.

> **EXAMPLE | Acme Billing**
>
> Acme had a capacity plan updated each year with sales forecasts and the month-end invoicing peak. **Rating: C.**

14. Performance and resource use monitored with thresholds

Monitoring shows whether resources are approaching their limits before users notice. Alert thresholds should trigger action early enough to prevent degradation.

Evidence to request

- Monitoring dashboards for CPU, memory, storage, database connections and response times.
- Defined alert thresholds and escalation routes.
- Examples of alerts that triggered action before an incident.

Red flags

- Alerts fire so often they are ignored.
- Monitoring covers infrastructure but not end-user response times.

> **EXAMPLE | Acme Billing**
>
> Infrastructure metrics and end-user response times are monitored with tiered alerts. **Rating: C.**

15. Load and scalability testing covers peak events

Load testing shows how the service behaves under peak demand. It should reflect realistic peaks such as month-end processing, sales events or product launches, not just average traffic.

Evidence to request

- Load test plans and results for critical services.
- Test scenarios that match known peak events.
- Actions taken from test findings.

Red flags

- The last load test was run before a major architecture change.
- Tests only cover the web tier, not databases or third-party APIs.

> **EXAMPLE | Acme Billing**
>
> Load tests are run before each major release but do not include the payment provider's API limits. **Rating: P.**

16. Cloud quotas, service limits and auto-scaling reviewed

Cloud resources are governed by account quotas, regional service limits and auto-scaling settings. These hidden ceilings often cause capacity failures at the worst possible moment.

Evidence to request

- A record of relevant quotas and service limits for critical services.
- Auto-scaling configuration with minimum and maximum limits, and the reasoning behind them.
- Evidence of periodic review and quota increase requests where needed.

Red flags

- Maximum auto-scaling limits were set during initial set-up and never revisited.
- Nobody knows the account-level quota for key resources.

> **EXAMPLE | Acme Billing**
>
> Auto-scaling limits are documented, but account quotas have never been reviewed. **Rating: P.**

17. Capacity risks assessed and recorded

Known capacity constraints, such as a database that cannot scale horizontally or a licence limit, should be recorded as risks with owners and treatment plans.

Evidence to request

- Capacity-related entries in the risk register.
- Treatment plans with owners and dates.

Red flags

- Capacity issues are known to engineers but not recorded anywhere.

EXAMPLE | Acme Billing

The single-region database limit was known to engineering but never recorded as a risk.

Rating: G.

18. Recurring performance problems analysed and acted on

Repeated slowdowns or near-misses are early warnings. They should be analysed through problem management and lead to lasting fixes.

Evidence to request

- Trend reports on performance incidents.
- Problem records for recurring performance issues, with root cause and fixes.

Red flags

- The same slowdown is "fixed" by restarting a service every month.

EXAMPLE | Acme Billing

Month-end slowdowns led to a problem record, root cause analysis and a database index fix.

Rating: C.

GOOD EVIDENCE | What good looks like for Capacity Management

- A capacity plan linked to business forecasts and reviewed regularly.
- Monitoring and alerting that catches problems before users do.
- Load tests that reflect real peak events, including third-party dependencies.
- Documented quotas, limits and capacity risks with owners.

Process 3 rating table

#	Check	Rating	Evidence reviewed	Notes
13	Capacity requirements and forecasts exist	☐C ☐P ☐G ☐N/A		
14	Performance and resource use monitored with thresholds	☐C ☐P ☐G ☐N/A		

15	Load and scalability testing covers peak events	☐C ☐P ☐G ☐N/A		
16	Cloud quotas, service limits and auto-scaling reviewed	☐C ☐P ☐G ☐N/A		
17	Capacity risks assessed and recorded	☐C ☐P ☐G ☐N/A		
18	Recurring performance problems analysed and acted on	☐C ☐P ☐G ☐N/A		

Process 4: Incident Management and Postmortem Reviews

ISO/IEC 20000-1:2018 Clauses 8.6.1, 8.6.3 and 10.1 • Checks 19–24

Incident management restores normal service as quickly as possible and limits the impact on the business. Problem management (8.6.3) and corrective action (10.1) make sure the organisation learns from incidents and prevents them from happening again.

The incident postmortem connects these three: it explains what happened, why, how the organisation responded and what will change.

> **OUTAGE LENS | The response matters as much as the root cause**
>
> Acme Billing did not cause the cloud outage, but it fully owned its response: how fast it detected the issue, who it escalated to, what it told customers and what it learned. That is what the auditor evaluates.

19. Incidents identified, classified and prioritised by business impact

Incidents should be detected quickly and classified according to their impact on the business, not just on technology. A cloud outage affecting payments should be recognised as a major incident immediately.

Evidence to request

- Incident classification and priority matrix based on business impact and urgency.
- Incident records showing detection time, classification and priority.
- Monitoring that detects service-level impact, not only infrastructure failures.

Red flags

- The first notification of the outage came from customers rather than monitoring.
- Major incidents are classified the same way as routine tickets.

> **EXAMPLE | Acme Billing**
>
> Monitoring alerted within 4 minutes, and the incident was classified as Priority 1 within 10 minutes. **Rating: C.**

20. Major-incident escalation paths and roles defined and used

A major incident needs clear roles: an incident manager, technical leads, a communications owner and an executive contact. Escalation paths, including to suppliers, should be defined and actually followed.

Evidence to request

- Major incident procedure with defined roles and escalation paths.
- Incident records showing who was involved and when.
- Evidence of escalation to the supplier.

Red flags

- Engineers worked the issue for an hour before anyone told management.
- There was no route to escalate to the supplier.

> **EXAMPLE | Acme Billing**
>
> The major incident procedure was followed internally, but there was no supplier escalation route (see check 3). **Rating: P.**

21. Stakeholders and customers updated during disruptions

Customers and internal stakeholders need regular, accurate updates during a major incident. Silence damages trust more than the outage itself.

Evidence to request

- Communication procedure with templates, channels and update frequency.
- Status page history or customer communications sent during the incident.

- Internal updates to leadership, support and account management teams.

Red flags

- The status page still showed "All systems operational" during the outage.
- Support staff learned about the outage from customer calls.

> **EXAMPLE | Acme Billing**
>
> The first customer update went out 75 minutes after detection, and the status page was not updated until the second hour. Support received over 600 tickets. **Rating: G.**

22. Postmortems cover detection, escalation, communication and root cause

A useful postmortem goes well beyond recording the outage duration. It should reconstruct the timeline, assess how well each part of the response worked, identify contributing factors on the organisation's side and recommend changes.

Evidence to request

- Postmortem reports for major incidents.
- A timeline from detection to full recovery.
- Analysis of detection, escalation, communication, recovery and contributing factors.
- Agreed actions with owners.

Red flags

- The postmortem reads "external cause, no action required".
- Postmortems blame individuals rather than examining processes.

> **EXAMPLE | Acme Billing**
>
> The initial incident report read only "Cloud provider outage, 4h 07m, resolved by provider".
>
> After the audit query, a full postmortem was started. **Rating: P.**

23. Corrective actions owned, dated and tracked to closure

ISO 20000 clause 10.1 requires the organisation to take action on nonconformities and evaluate whether those actions were effective. Postmortem actions are only valuable if they are completed.

Evidence to request

- A corrective action log with owners, due dates and status.
- Evidence of closure and effectiveness review.
- Overdue actions escalated to management.

Red flags

- Actions from earlier incidents are still open months later.
- Actions are closed without any check that they worked.

> **EXAMPLE | Acme Billing**
>
> Corrective actions are logged in the ITSM tool with owners and due dates, and are reviewed in the monthly service review. **Rating: C.**

24. Recurring incidents linked to problem management

When similar incidents recur, problem management should identify the underlying cause and drive a permanent fix or workaround.

Evidence to request

- Problem records linked to related incidents.

- Trend analysis that identifies recurring incident types.
- Known error records and workarounds.

Red flags

- Each outage is treated as a separate, unrelated event.

EXAMPLE | Acme Billing

Acme's ITSM tool links incidents to problems, and the monthly trend review raised a problem for supplier-related outages. **Rating: C.**

TIP | Mini postmortem quality test: 5 questions for any incident report

Ask these five yes/no questions of any major incident report. Three or more "no" answers mean the postmortem is not giving the organisation real learning:

- Does it include a **timeline** from first detection to full recovery?
- Does it assess **how quickly** the issue was detected and escalated?
- Does it evaluate **communication** with customers and stakeholders?
- Does it identify **contributing factors on our side**, even when the trigger was external?

- Does it produce **corrective actions** with owners and dates?

GOOD EVIDENCE | What good looks like for Incident Management & Postmortems

- Business-impact-based priorities and fast, monitoring-led detection.
- Defined major incident roles, including supplier escalation.
- Timely, regular customer communication through agreed channels.
- Blameless postmortems with actions tracked to closure.

Process 4 rating table

#	Check	Rating	Evidence reviewed	Notes

19	Incidents identified, classified and prioritised by business impact	☐C ☐P ☐G ☐N/A		
20	Major-incident escalation paths and roles defined and used	☐C ☐P ☐G ☐N/A		
21	Stakeholders and customers updated during disruptions	☐C ☐P ☐G ☐N/A		
22	Postmortems cover detection, escalation, communication and root cause	☐C ☐P ☐G ☐N/A		
23	Corrective actions owned, dated and tracked to closure	☐C ☐P ☐G ☐N/A		
24	Recurring incidents linked to problem management	☐C ☐P ☐G ☐N/A		

Process 5: Service Level Management

ISO/IEC 20000-1:2018 Clause 8.3.3 • Checks 25–30

Service level management agrees service targets with customers, monitors performance against them and acts when they are missed. For cloud-dependent services, the organisation's customer SLAs must be realistic given what its suppliers actually commit to.

> **OUTAGE LENS | The cloud SLA can be met while your service is down**
>
> A provider can meet its monthly availability commitment even after a 4-hour regional outage, while your customer-facing service misses its own target badly. Service level management checks whether that gap was understood and managed.

25. SLAs reflect business requirements

Customer SLAs should be based on what the business and customers need, such as availability during business hours or payment processing times, not simply copied from what suppliers offer.

Evidence to request

- Customer SLAs and the business requirements they are based on.
- Records of SLA agreement with customers or business owners.

Red flags

- SLA targets were copied from the cloud provider's terms.
- No customer or business owner has signed off the SLAs.

> **EXAMPLE | Acme Billing**
>
> Acme's SLAs were agreed with customers and reflect month-end invoicing needs. **Rating: C.**

26. Availability targets defined end to end

Availability should be measured for the full service the customer uses, including the application, database, identity, network and third-party APIs, not just the underlying infrastructure.

Evidence to request

- End-to-end availability definitions for customer-facing services.
- Synthetic monitoring or user-journey checks that measure availability as customers experience it.

Red flags

- Availability is measured only as server or virtual machine uptime.

> **EXAMPLE | Acme Billing**
>
> Availability is measured by synthetic login and payment transactions every minute. **Rating: C.**

27. Supplier commitments mapped to internal SLAs, with gaps recorded

If a customer SLA promises more than the underlying suppliers commit to, the gap must be covered by architecture (for example, redundancy) or recorded as an accepted risk.

Evidence to request

- A mapping of customer SLAs to supporting supplier commitments.
- Identified gaps with a treatment decision (redundancy, risk acceptance or SLA change).

Red flags

- The customer SLA promises higher availability than the single cloud region it runs in can support.

> **EXAMPLE | Acme Billing**
>
> Acme promised 99.9% monthly availability to customers but had never compared it with the provider's commitment for a single-region deployment. **Rating: G.**

28. Service performance reported and reviewed with customers

Regular service reports show customers and management how services are performing against SLAs and highlight trends before they become breaches.

Evidence to request

- Monthly or quarterly service reports.
- Records of service review meetings with customers or business owners.

Red flags

- Reports are produced but never discussed.

> **EXAMPLE | Acme Billing**
>
> Monthly reports are sent to key customers, but review meetings only happen for the top 10 accounts. **Rating: P.**

29. SLA breaches lead to corrective actions

A missed SLA should trigger analysis and action, such as service improvement, architecture change, supplier discussion or SLA renegotiation, not only service credits.

Evidence to request

- Records of SLA breaches and the actions taken.
- Service improvement plans linked to breaches.

Red flags

- Breaches are settled with service credits and then forgotten.

> **EXAMPLE | Acme Billing**
>
> The outage breached the monthly SLA. Service credits were issued, but no improvement action was raised against the SLA itself. **Rating: P.**

30. Customers informed of significant disruptions within agreed timeframes

SLAs or customer agreements should define how and when customers are notified of significant disruptions, and the organisation should meet those commitments.

Evidence to request

- Notification commitments in customer agreements.
- Evidence of notification times for recent major incidents.

Red flags

- The agreement promises notification within 30 minutes, but no one tracks whether this is met.

EXAMPLE | Acme Billing

The agreement promised notification within 30 minutes; the first notice went out after 75 minutes. **Rating: G.**

GOOD EVIDENCE | What good looks like for Service Level Management

- SLAs agreed with customers and based on business needs.

- End-to-end availability measured the way customers experience it.
- A documented map from customer SLAs to supplier commitments, with gaps treated.
- Breaches that lead to real improvement, not just service credits.

Process 5 rating table

#	Check	Rating	Evidence reviewed	Notes
25	SLAs reflect business requirements	☐C ☐P ☐G ☐N/A		
26	Availability targets defined end to end	☐C ☐P ☐G ☐N/A		
27	Supplier commitments mapped to internal SLAs, gaps recorded	☐C ☐P ☐G ☐N/A		
28	Service performance reported and reviewed with customers	☐C ☐P ☐G ☐N/A		

29	SLA breaches lead to corrective actions	☐C ☐P ☐G ☐N/A		
30	Customers informed of significant disruptions on time	☐C ☐P ☐G ☐N/A		

Gap Summary and Action Plan

Once all 30 checks are rated, total the results for each process to see where the service management system is strongest and where it is most exposed.

Summary scorecard

Process	C	P	G	Readiness %
Supplier Management				
Service Continuity				
Capacity Management				

Incident & Postmortem				
Service Level Management				
Overall				

How to calculate readiness

Readiness % = (C + 0.5 × P) ÷ number of applicable checks × 100

Partial ratings count as half, because the process exists but has not been proven to work. N/A checks are left out of the count.

Worked example: Acme Billing after the outage

Process	C	P	G	Readiness %
Supplier Management	1	2	3	33%
Service Continuity	0	4	2	33%

Capacity Management	3	2	1	67%
Incident & Postmortem	3	2	1	67%
Service Level Management	2	2	2	50%
Overall	**9**	**12**	**9**	**50%**

WATCH OUT | Acme Billing: 50% overall, with significant gaps

Acme's capacity and incident management are in reasonable shape, but **service continuity**, **supplier management** and **service level management** are where the outage did the most damage, and where a certification body would most likely raise nonconformities.

The pattern is typical: the organisation was good at responding to incidents it caused itself, but weak at planning for a supplier failure it could not control.

How to read your result

Readiness	Status	What it means

80% or more	Audit-ready	Processes are working with evidence. Focus on closing remaining Partial items and keeping records current.
60–79%	Fix before the audit	Core processes exist but key evidence is missing. Prioritise high-risk gaps and gather evidence over the next audit cycle.
Below 60%	Significant gaps	Major nonconformities are likely. Build a structured remediation plan before booking a certification or surveillance audit.

How to prioritise the gaps

Not every gap carries the same risk. Fix gaps in this order:

1. **Gaps that affect customer-facing critical services**, especially those with no workaround (for example, no continuity plan for the payment API).
2. **Gaps that would be a major nonconformity**, where a required process is missing entirely.
3. **Gaps exposed by a real incident**, because the risk has already materialised once.

4. **Partial ratings that only need evidence**, which are often quick wins (for example, recording supplier review minutes).

Action plan template

Gap #	Finding	Risk	Corrective action	Owner	Due date	Status

Example: Acme Billing action plan (first three entries)

Gap #	Finding	Risk	Corrective action	Owner	Due date	Status
9	No continuity plan for cloud region failure	H	Add region-failure scenario, invocation criteria and roles to the plan; test within 90 days	Head of IT Ops	Q1	Open
21	First customer update 75 min after detection	H	Pre-approved status templates; communications owner in P1 roster; 30-minute update rule	Support Lead	30 days	In progress
27	Customer SLA not mapped to supplier commitment	H	Map SLAs to supplier terms; decide on multi-region design or revise SLA	Service Manager	60 days	Open

Bonus: tracing an outage to audit evidence in 5 steps

Use a real outage as an audit case study. Instead of reviewing procedures in isolation, follow one incident from detection to recovery:

Step	What to do	Checks it covers
1. Understand the incident	Review the postmortem, timeline, affected services and business impact.	19, 22
2. Identify relevant processes	Map the incident to supplier, continuity, capacity, incident and service level management.	All five processes
3. Examine supporting evidence	Request contracts, monitoring reports, incident records, continuity test results and corrective action logs.	3, 4, 11, 14, 23
4. Identify process weaknesses	Find missing controls, unclear responsibilities or procedures that did not work as written.	5, 9, 20, 27
5. Verify corrective actions	Confirm that weaknesses were addressed and improvements were tested.	6, 12, 23, 29

NEXT STEP

Audit Cloud-Era IT Services With Confidence

This checklist helps you find gaps in one service management system. The **Certified ISO 20000:2018 Lead Auditor** certification from GSDC helps you plan and lead complete ISO 20000 audits, and tell the difference between a system that is only documented and one that truly works.

Build practical skills in:

- **Audit planning:** scoping and planning ISO 20000 audits for cloud-dependent services.
- **Evidence collection:** gathering and evaluating objective audit evidence.
- **Process evaluation:** assessing whether service management processes are implemented and effective.
- **Nonconformity identification:** recognising, grading and reporting nonconformities.

Explore the Certified ISO 20000:2018 Lead Auditor Certification

[ENROLLMENT LINK]

CERTIFIED ISO 20000:2018 LEAD AUDITOR

LEARN DIRECTLY FROM GLOBAL PRACTITIONERS, IT SERVICE MANAGEMENT EXPERTS, AND INDUSTRY LEADERS WHO ARE SHAPING THE FUTURE OF ITSM GOVERNANCE, SERVICE EXCELLENCE, AND MANAGEMENT SYSTEM AUDITING.



ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

LEARNING OBJECTIVE

- Understand risk-based auditing approaches within IT service management environments.
- Increase your skills in leading the audit team and managing the audit programs.

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org