QOS MSP
Managed Services

INFRASTRUCTURE MONITORING & MAINTENANCE:

A PRACTICAL CHECKLIST FOR BUSINESS OWNERS

+1 (877) 800-7672

www.qosmsp.com

Infrastructure Monitoring & Maintenance: A Practical Checklist for Business Owners

Most infrastructure failures are not mysteries. They are missed warnings — a disk that filled up, a patch that never got applied, a backup that quietly stopped running weeks before anyone needed it.

The businesses that avoid outages are not the ones with the newest hardware. They are the ones where someone is watching the right signals, on a schedule, before those signals turn into downtime.

Use this checklist to see what should be monitored, how often systems should be patched, when to plan for more capacity, and how to prove your disaster recovery actually works.

1. Monitor These Every Day

Daily monitoring is where outages get caught while they are still small problems. These signals should be watched automatically, around the clock — not checked by hand when someone remembers.

Watch Daily:

- **Server availability** — every critical server responding and online
- **Disk space** — free capacity on every volume, with alerts before it fills
- **CPU and memory load** — sustained spikes that signal a failing or overloaded system
- **Backup jobs** — last night's backups completed and reported success
- **Critical services** — databases, email, line-of-business apps actually running
- **Security alerts** — failed logins, new admin accounts, unexpected changes

The Rule:

If a signal can take your business offline, it belongs on the daily list — monitored and alerting, not spot-checked.

2. Monitor These Every Week

Weekly checks catch the slow-moving problems — the trends that never trip a daily alarm but end in failure if nobody is looking.

Review Weekly:

- Disk, memory, and bandwidth **trends** over the past week
- Server event and error logs for repeating warnings
- Hardware health — drive SMART status, RAID array state, power supplies
- Patch status — what is pending, what failed to apply
- Backup storage capacity and retention headroom
- Certificate and license expiration dates coming up

The Red Flag:

A metric that has climbed every week for a month will not stop climbing on its own. Trend lines are early warnings — act on them.

3. Patch on a Cadence, Not on Panic

Unpatched systems are the most common way businesses get breached and the most avoidable. The goal is a predictable schedule, so patching happens before an incident forces it.

A Practical Patch Cadence:

- **Critical security patches** — tested and deployed within **72 hours** of release
- **Operating system updates** — a scheduled **monthly** patch window per environment
- **Application and browser updates** — monthly, or automatically where safe
- **Firmware** (servers, switches, firewalls) — reviewed **quarterly**
- **Emergency (zero-day) patches** — out-of-cycle, as soon as tested

The Goal:

Test first, deploy in a window, keep a rollback plan. Patching should be routine and boring — never a fire drill.

4. Watch These Capacity Triggers

Capacity problems are predictable if you set a line and act when you cross it. These are the thresholds that mean “plan to scale now” — not “wait until it breaks.”

Plan to Scale When:

- **Disk usage** passes **80%** on any critical volume
- **Memory** runs sustained above **85%** during normal hours
- **CPU** holds above **75%** at routine peak load
- **Network bandwidth** regularly tops **70%** of the link at peak
- **Backup windows** start running past their scheduled time
- **User or device count** is projected to grow 20% or more this year

The Rule:

Capacity should be added on a plan with lead time, not bought overnight at a premium after something fails.

5. Verify Backups — Don’t Assume Them

A backup that has never been restored is a hope, not a safety net. Most businesses discover their backups were failing only when they try to recover.

Confirm Every Backup:

- Runs on schedule and **reports success** — monitored, not assumed
- Follows a **3-2-1 approach** — three copies, two media, one off-site
- Covers every critical system — not just file servers
- Is **encrypted** and protected against ransomware and deletion
- Meets a defined **retention** period the business has agreed to

The Red Flag:

If nobody can name the date of the last successful test restore, you do not have backups you can rely on.

6. Test Disaster Recovery on a Schedule

Disaster recovery is only real once it has been tested. A recovery plan on paper tells you what should happen; a test tells you what actually will.

A Practical DR-Test Schedule:

- **Monthly** — restore a sample of files and one system to confirm backups work
- **Quarterly** — recover a critical server or application end to end
- **Semi-annually** — a tabletop walkthrough of the full recovery plan with the team
- **Annually** — a full failover test against your recovery-time objective

The Goal:

Know your **recovery time** and **recovery point** objectives — how fast you must be back and how much data you can afford to lose — and prove you can meet them.

7. Track Hardware and Vendor Lifecycles

Equipment and software do not fail on your schedule. Tracking their lifecycles turns end-of-life surprises into planned, budgeted replacements.

Keep a Current Record Of:

- Age and warranty status of every server, switch, and firewall
- **End-of-life and end-of-support** dates for hardware and operating systems
- Software and license renewal dates
- Vendor and support contacts for each critical system
- A rolling 12-to-36-month replacement and budget plan

The Bottom Line:

Hardware running past end of support is not saving money — it is deferring a bill that grows with every month of risk.

8. Report on Health Every Month

Monitoring only pays off when someone reviews it. A monthly report turns raw data into decisions and keeps infrastructure from drifting out of sight.

A Good Monthly Report Shows:

- Uptime and any incidents, with what caused them
- Patch compliance — what is current and what is behind
- Capacity trends and anything approaching a threshold
- Backup success rate and test-restore results
- Aging hardware and upcoming renewals
- Recommendations, prioritized by risk

The Goal:

An infrastructure you can see — where nothing important is a surprise and every decision has data behind it.

The Final Word

Stop treating infrastructure as something you only think about when it breaks, and start treating it as a system that is watched, maintained, and reported on every month.

Businesses that monitor the right signals, patch on a cadence, plan capacity ahead, and test recovery avoid the outages that catch everyone else off guard. The rest find out what was wrong the hard way.

Next Step:

Walk this checklist against your own environment. For each item, ask one question: do we do this, on a schedule, and can we prove it? Every “no” is a gap worth closing before it closes on you. If you would rather one accountable team owned the whole list, that is exactly what infrastructure management from QOS MSP does.