

3 EASY TIPS TO CUT DOWNTIME BY 50%

Maximize Performance and Availability While Lowering Cost

In today's world, we need to keep our organizations up and running 24/7 with limited or non-existent unscheduled downtime, or face the wrath of IT directors and CIOs. The key is to apply the principles of effective server monitoring and to understand the pitfalls that continue to take down systems today. By understanding how to manage CPU and Resource overload, Disk Space, and IT Services in your environment, you can help cut unscheduled downtime by 50% or more. To find out how, read on.

TIP 1 IT SERVICES

One of the most frequent causes of unscheduled downtime for critical applications is a critical service stopping or stalling. In fact, a stalled service can be more disruptive than an outage due to the difficulty in detecting it.

In this example, let's look at a Windows environment with three critical applications that need their services monitored: SQL Server, IIS, and Exchange.

While all three are typically used for mission-critical applications, Exchange can be the most susceptible to a service stopping or stalling. The results can be painful, and it happens far too often. A four-hour time span of unavailability for a major enterprises email system can be extremely expensive, and cause executive management to start questioning the competency of the entire IT department. In this example, email is taking longer and longer to be delivered. As all IT staff know, business users instantly notice when their email is not functioning at 100%, and they start flooding the help desk with calls.

Two Windows services properties are important to note in this example: Windows services are usually:

1. Critical and 2. Hidden. For example, consider the spooler service. The Service Control Manager's applet interface (the icon of two intermeshed straight-cut gears) can be used to stop the spooler service. However, if you try to print you will receive an error message about printers not being installed. Although this is just one small microcosm, and doesn't sound as mission-critical as losing Exchange, imagine the impact of something similar across infrastructure and business units including HR, Operations, and Marketing. In a few extreme cases, the production server will need to be rebooted. Fortunately, this is very rare.

Real Life Example 1 - A Large US Credit Union with over 50 Years in business and more than \$1 Billion in assets.

Originally focused on purchasing a simple event log monitoring product, this Credit Union quickly saw the overall benefit in monitoring all aspects of their Linux, Windows and VMware infrastructure. The key was managing these systems from a single IT dashboard. Uptime Infrastructure Monitor is now used in this environment to monitor various systems, including Linux and Windows servers for performance and availability metrics on Exchange and SQL Server, as well as deep reporting on VMware performance and capacity.

In addition, due to management pressure and IT Service Level Agreement (SLA) initiatives, Uptime Infrastructure Monitor's SLA monitoring and management features are being used to report on server uptime, CPU utilization, and more. In their search for a Service Level Management tool, this client considered purchasing multiple products to meet their service level requirement needs. However, Uptime Infrastructure Monitor addressed all their monitoring, alerting, and reporting needs (including SLAs) with a single, unified IT monitoring solution.

Today, this client monitors over 700 servers and network devices and raves about Uptime Infrastructure Monitor's ease of deployment and usability compared to some of the bigger frameworks, like HP Openview, IBM Tivoli, and BMC Patrol.

EXTRA TIP

Ensure your services are fully monitored. Send automated, scheduled reports to your team or management on a regular (weekly or monthly) basis.

TIP 2 CPU & RESOURCE OVERLOAD

Of the three major causes of critical application failure, this might seem to be the least severe, as an application can continue to limp along during times of CPU and resource overload. However, this is exactly what makes it so dangerous.

Why? This type of issue can fester over months without being noticed: nothing is broken, so nothing is fixed. What happens? The email systems do not simply stop delivering email, as often happens with a service(s) stalling. The danger with resource and CPU problems is that email just takes longer and longer to be delivered. What do you think is the first application that end users tend to complain about? Exactly: email. Once again, managers start to blame the IT department and escalate those complaints. Eventually the complaints filter upward to the CEO, who has noticed that email delivery has been slow as well. Guess who gets the blame? And guess who's SLA is now come into question? All this pain for a problem that should have been deleted and fixed before end-users noticed.

The underlying truth is that CPU and resource overload can have a seriously adverse impact on application efficiency, and especially on mission-critical applications.

Let's look at an example with SQL Server. We'll assume that SQL queries are taking increasingly longer to complete, and the result is lower end-user productivity and higher end-user frustration. If there are 500 end users, and the typical query takes 15% longer, the lost productivity quickly adds up. It's easy to see that the price of an upgraded server is often justified through one day's lost productivity. However, you need monitoring that can find the problem proactively.

So now, you know that more hardware is needed to solve the problem. However, how do you justify this request for additional budget to management? You need reporting that can provide management with quantifiable proof before they approve the budget (this type of monitoring and reporting is included in Uptime Infrastructure Monitor).

Two proactive IT management practices prevent these types of critical bottlenecks. The first is quite simple, requiring active monitoring of CPU processing and real-time alerting when thresholds are breached. The second includes leveraging historical performance data with deep reporting that can graph CPU growth. With this trend-line based graph and report, it becomes clear to management that new hardware is needed to solve the problem. If your software can't produce these types of historical performance reports with the detailed metrics and graphs you need, then you can't justify the request to management. Uptime Infrastructure Monitor is designed for the deep monitoring, alerting and reporting that mid and large enterprise companies require to handle these types of situations. So if you have Uptime Infrastructure Monitor, problem solved. Now, just sit back, enjoy the day, and wait for your promotion.

EXTRA TIP

Ensure that alert thresholds are set strategically and that you have access to the historical performance data of key metrics. Create trend-line analysis graphs to show key metric growth over day, week, month, or year. Nothing has the impact that a trend-line growth graph clearly showing management how close you are to running out of CPU or other resources, along with the consequences of that scenario.

Real Life Example 2 - A New York based Global Investment Bank with over \$1.4 trillion in assets & operations in 50+ countries.

At this large, global investment bank, a key requirement is accurately monitoring and reporting on mission-critical server resources, such as CPU, I/O, and file system capacity. It's critical that these, and other key metrics, remain within the defined thresholds specified by IT management. If these thresholds are exceeded, it requires immediate action on the part of the IT team. Their process of evaluating monitoring solutions included grueling head-to-head comparisons. In the end, they chose Uptime Infrastructure Monitor.

The client now has a unified IT dashboard with performance reports automatically emailed and posted to Web sites, allowing more than 50 IT analysts to monitor and report on over 900 UNIX, Windows, and VMware servers in New York and London, on a daily, weekly, and monthly basis. This constant and in-depth communication ensured that CPU, I/O, and file system capacity remained inside the defined thresholds and ran at optimal capacity. These reports highlighted server resource performance over the past hour, day, week, and month, so IT management had the opportunity to take pro-active measures before any potential crisis were reached.

Real Life Example 3 - A Leading Wireless Phone Company with over 20 Million Customers and \$8 Billion in Revenue

One of Canada's most respected wireless providers monitors 680 mission critical Solaris servers. The IT department wanted IT Service Level Agreement (SLA) monitoring and reporting to show the availability of key production servers, many running very large Oracle databases. At the same time, they were interested in proactive performance and capacity monitoring, alerting, and reporting to help maintain high service levels. The IT staff was stuck considering multiple products needing a complex integration... until they found Uptime Infrastructure Monitor.

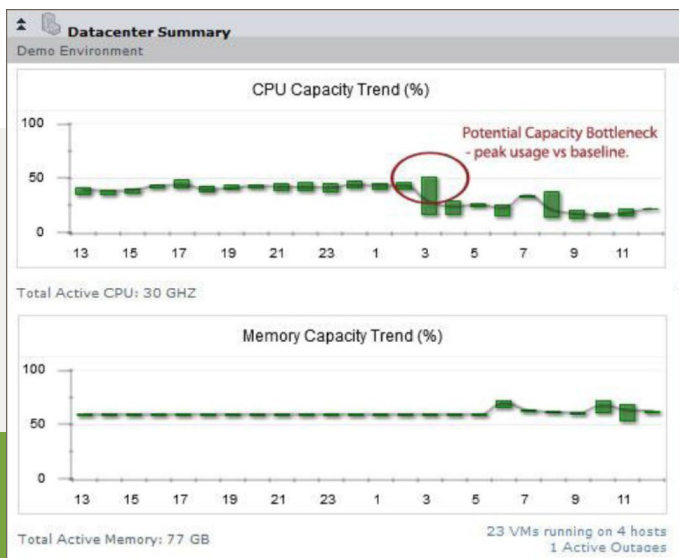
The provider chose Uptime Infrastructure Monitor for its comprehensive and reliable monitoring and alerting features for critical resource thresholds – including CPU, Memory, Network I/O and File System Capacity. Uptime Infrastructure Monitor also provided in-depth management and service-level reporting that ensured in-house capacity planners had the right information to make strategic capacity-related decisions.

TIP 3 DISK SPACE AND THE NEED TO PLAN FOR IT.

We saved the number one reason for unscheduled downtime for last, and we know you might be thinking, “Are you kidding? Of course I already know that.” You’re right, yet it still happens far too often, even in good IT departments. Why? Read on and find out.

Even though it might be obvious that disk space needs to be managed, it is essential to plan for disk space in light of the many other tasks in your infrastructure. When your environment is running at status quo, disk space is easily managed. However, when unexpected events happen, you don’t want to be caught with your pants down.

It is important to make certain your servers have alerting functions set to notify the right people of disk issues. Too much alerting can cause a similar effect to ‘crying wolf,’ so alerts need to be sent to those on a must-know list first. If an alert needs to be escalated, it should be escalated from there. Alerting best practices include times when disk space thresholds are approaching, not at, critical.



Datacenter, Cluster,
Resource Pool & vApp
Aware Monitoring

From here, your alerting function should send its notifications out via email, PDA, phone, SMS, etc. False or nonexistent alerts are par for the course with many IT monitoring tools, and that is unacceptable. Make sure your alerting solution is trust-worthy, effective, and consistent.

Lastly, remember to leave enough time to take action against these potentially critical disk issues. If the alert goes unanswered by the first recipient, ensure additional alerts are scheduled and sent as a failsafe.

EXTRA TIP

Disk Space should be on every IT manager or system administrator's expertise list, but it continues to be a major cause of outages. Make sure your monitoring software has a robust alerting system that is topology and dependency aware to keep you and your team in control, no matter what.

Real Life Example 4 - A Leading California Food Company with 250+ Warehouse Stores & \$500 Million in Annual Revenue.

Originally purchasing Uptime Infrastructure Monitor to consolidate AIX servers, increase server performance and capacity plan for additional resources, this company found the extra benefits of Uptime Infrastructure Monitor for its alerting and notifications where other products had failed.

An early detection of an unscheduled outage by Uptime Infrastructure Monitor saved this client over \$50,000 in one instance. The scary fact is that this same outage was completely missed by a large, expensive, and well-known Big 4 system management framework product.

The client explains how they use Uptime Infrastructure Monitor, "If something goes wrong that I'm responsible for (servers, infrastructure, and applications), I know about it before it becomes an issue. I can be more proactive in pinning down the problem. The historical trending is great, as it shows me whether the problem happened quickly or slowly over time."

He goes on to say, "Using Uptime Infrastructure Monitor, we were able to identify servers that were sub-optimally configured over the past two years. Uptime Infrastructure Monitor helped us find those problems and set the servers to perform at peak efficiency. The new configuration resulted in improved system performance. We showed that improvement to management with Uptime Infrastructure Monitor's reports, and they were very impressed with our work."

[For the full case study of this company, click here.](#)

“

With Uptime Infrastructure Monitor,
if something goes wrong that I'm responsible for,
I know about it before it becomes an issue.

”

FREE PRODUCT EVALUATION CHECKLIST & CALCULATOR

If you are considering evaluating unified server monitoring solutions, this **IT Systems Monitoring Checklist and Calculator** is an excellent way to start. It's designed to be vendor agnostic and customizable to help you compare different products. A free download is available here:

[Download Here](#)

