

Remote Analyst for HPE Integrity NonStop servers

Reduce costs and prevent unplanned outages



Contents

Is faster always better?	3
Design matters.....	3
Present and clear risks	3
Problems are not detected in time.....	3
Personal knowledge base dependency cannot be carried forward.....	4
Where the money goes.....	4
Outage.....	4
Human resources	4
Today's technology can help	4
Viable solution: Remote Analyst.....	5
Resource utilization levels	6
Load balancing analyses.....	6
Recommendations and expected results	7
Alerts and links to details	7
Better or worse than before.....	7
Multiple years of trend analyses.....	7
Summary to detail drill-down options.....	8
Deep analyses for all resources	8
What-if analyses	9
Memory leak detection.....	9
Storage capacity analyses and predictions.....	10
Detail network analyses.....	11
The solution is available now	11
Early warning/outage avoidance.....	11
Reduced cost.....	11
Less management requirements	11
Use our knowledge to your advantage.....	12



Is faster always better?

These days computers are significantly more powerful than ever before. A remarkable benefit is that more computing can be done in a smaller time/space while using less processing resources (such as CPUs). One result of these benefits that can be overlooked is the effect upon possible application outages.

Factors that affect system performance are now buried deeper in the OS and often go unnoticed by today's online performance monitors. CPU busy, memory usage, queues, and many other surface-level metrics are often watched by these performance tools. The tools may show that there are no issues, while a problem lurking deep in the system can go unnoticed. To use a human analogy, all surface vital signs may be fine, while an undetected disease develops into a serious issue.

Business activities can be in trouble by the time performance issues surface and are displayed by an online monitor. Earlier problems may have already spread into other areas, an outage has taken place, or is imminent. In those cases, it can be costly and time-consuming to recover from an outage or to locate the root cause.

To be clear, we are not referring to hardware, network, or even OS software failures. On HPE Integrity NonStop servers, these are already addressed as part of the inherent fault-tolerant architecture. The type of issues affecting customers today are often due to improper resource utilization, weak points in application design, or resource thresholds, which once reached could affect many other areas. These are just a few possibilities. There are many others, and nearly all combinations are possible, as one issue can easily trigger another. The problem is made worse as you try to determine the root cause while bypassing noisy symptoms.

Another aspect of faster computing environments is that people are no longer able to keep up with effective monitoring and analyses of their systems' performance. This can happen regardless of the number of displays or screens, the number of people involved, or their level of expertise. Although many current performance experts have decades of experience, some continue to use outdated trusted tools. In today's computing environment, trusted tools from the past may no longer be effective. The tools may use obsolete technologies and also rely upon an expansive knowledge base by system operators. This could be a Catch-22 scenario, where people rely on tools and the tools rely upon people for their expert knowledge.

Nearly all of today's surface-level monitors perform reporting while leaving people to detect issues, analyze possible options, and then decide what to do. Another important aspect is that the next-generation of analysts may be unable or unwilling to work with these tools.

HPE NonStop servers are best at providing their customers with a fast and uninterrupted transaction environment. Whether from a software or hardware point of view, every aspect of an HPE NonStop server is built with enterprise-level quality in mind. Many of today's online HPE NonStop performance monitoring tools use storage and computing resources of an HPE NonStop server to do their reporting and analyses. What if you could improve upon the performance monitoring, reporting, and forecasting ability and do it without consuming valuable system resources? That is part of what this white paper intends to share with you.

Design matters

Many of the online monitoring tools employed today were originally developed over a decade ago. They were bound to the available technology back then when internet and off-production computing was not an available option. Over the years, these tools grew in size, offering more features, screens, and reports, yet staying fixed to their original design framework and limitations. What this has meant to their users is an elevated level of risk and increased cost.

Present and clear risks

Aside from the major and present possibilities of a system outage, there are two significant risk factors:

Problems are not detected in time

- Online screens that monitor high-level resource usage (for example, CPU, process, top-level queues, CPU/memory, network) are unable to look deep to detect low-level contentions or imbalances. There is too much data for them to analyze promptly. Analyses of this data can take a substantial amount of computing work on your server, which can impact valuable business applications resources.
- In addition, some tools don't have time to understand the overall behavior of the system and are busy showing what is currently happening. By the time an online monitor notices an issue, it can be already too late, customers are affected, and the recovery will be time-consuming. At best, there could be a reversion or restart of the same configuration, which eventually leads to the same issues. These types of problems should have been detected weeks in advance along with recommended solutions that allow plenty of time to implement a fix.



Personal knowledge base dependency cannot be carried forward

- Expert employees who have spent a large part of their career on a given system/platform are less available in the market today. That situation is further complicated by a newer generation of employees who do not have the years of specific experience that the prior experts have built.
- Long-term training and the attainment of an expert knowledge base is a declining practice. In addition, having a new team member use the old tools can be like turning back the clock to decades earlier and starting over. This is a risk element that no company can afford to take as we move forward with technology.

Where the money goes

Increased cost is reflected in several areas:

Outage

- An outage can occur due to improper configuration of a system, network, and/or application resources, rather than hardware or system software failure. When the desired response time of 10 milliseconds is increased to 10 seconds—an outage is felt by the person waiting for a transaction to complete. Imagine if a transaction drops when certain thresholds are reached, while the systems are fully operational. At a glance, all may seem well (for example, CPU busy is well below 50%), the system is up, but to anyone depending upon that transaction, they just had an outage.
- Outage cost is hard to determine, as it varies from one business to the next. However, in all instances, it is a major factor in terms of current lost business revenue and future lost opportunities. Bad publicity in the market worsens the situation and a company's reputation can easily be tarnished. This is a major and immeasurable cost.

Human resources

- Additional dependence on human analyses is costly, both short and long term. Over time, the cost of computing resources has gone down, while the resource costs of people have been on a steady increase. Furthermore, training, education requirements, and maintaining backup resources increases the cost.

Today's technology can help

Most of the available tools are based on the late 20th century technologies. Also, the most important technological achievements (needed for automated analytics) were not readily or practically available. As a result, performance tool designers had to limit their scope and set their features to navigate the limitations (from using keyboard characters to construct a graph, to Visual Basic implementations). Though commendable at the time, these designs don't meet today's needs and have no place in today's open environment.

Here is a list of top technologies that have significantly impacted the way we do business today:

- Some of the major technological innovations of the recent era have been availability, bandwidth, and most notably the speed of network traffic. These innovations have made it practical to move large volumes of performance metrics off the HPE NonStop server and to an environment better suited for analytic work.
- Secured File Transfer Protocol (SFTP) and Secure Shell (SSH) access further meet the audit requirements at sensitive installations. We can now transfer gigabytes of data off the HPE NonStop server with ease of mind.
- Encryption features and firewalls help us keep the metrics off production.
- Cloud computing offers inexpensive and elastic computing resources, which takes on all incoming load and scales up or down as needed.
- Inexpensive off-NonStop storage offers the option to maintain years of performance metrics for deep analytics on system and application behavior.
- Viable email and graphic communication options offer immediate reach to customers to warn of impending issues.

Put them all together, and you have the pathway to new and improved analytic methods. Take any of them away and you are faced with an impossible or impractical solution. Additionally, the optimal solution needs to leverage all the options that can also invent new designs to take the analyses even further.

First, gathering data needs to be offloaded from the HPE NonStop server, so that we do not impact production activities while performing business analytics. This requires high bandwidth and secure connections. Inexpensive, expansive, and reliable storage is required from the server. This storage contains large volumes of data to serve as a platform for behavior analyses and prediction algorithms.



Off-production computing resources for analytics need to be powerful, cost-effective, and flexible. Queuing of resources is not the best option, nor the requirement to hold a large set of acquired resources in reserve. We are referring to the need for elastic computing and storage that is available only when needed.

Last, but perhaps the most important piece is the ability to effectively establish communication with users. This requires sending comprehensive emails to different users, based on their desired criteria and schedules. The users also need a method where they can securely access the systems' metrics, review past analyses, request new ones, or manually do further analyses on their own. Customers need to do this from anywhere—a smartphone (iPhone® or Android), tablet, or laptop (Mac or PC). No client application is required. Moreover, this must be done in a secured environment with no connection to any of the production HPE NonStop servers.

From a technological point of view, this is a long list of requirements. It takes a significant amount of time for all pieces to come together. However, once all the pieces are in place, you still need to develop a software platform and architecture that takes advantage of the new possibilities.

One significant part of the new solution would be to allow off-site computers to do complete analyses of HPE NonStop servers. In short, we need an elastic computing base to keep pace with the ever faster and more complex resources of HPE NonStop servers and applications.

Even if humans are available 24x7, observing multiple monitors, they cannot keep pace with everything that happens on an HPE NonStop server.

It is no longer enough to show a problem on the screen when it happens. Issues must be detected in advance before they rise to the surface. Solutions must be recommended along with a selection of actions to implement well before any screen-based alarm is raised. This would require going through large volumes of data continuously throughout the day, every day, and at high computing speeds. In essence, it would require other computers to constantly monitor and analyze all the system and application activities of the HPE NonStop server.

Industry leaders already understand this and have solutions in place. Large online computing environments use their cloud-based elastic solutions to keep a tab of their main network of computers that manage their daily business. A second layer, and perhaps more interesting implementation, is online computing environments that employ cloud-based computing resources to monitor and manage the cloud business offered to customers worldwide. They employ the same architecture described here, the only difference being that they do not analyze HPE NonStop server.

There is only one fully automated, off-production analytics solution for HPE NonStop servers. All other analytics products use the server to perform their analyses.

Using flood conditions as an example, nearly all of these solutions are best at providing surface monitoring and providing alarms when flood levels reach the front door. It helps to repeat that this analogy requires us to have a solution in place where we design and build dams, build water redistribution channels, put in the required resources, and run drills long before the first raindrop falls.

Viable solution: Remote Analyst

Remote Analyst is a fully automated performance analytic service for HPE NonStop servers.

As you are reading this document, gigabytes of encrypted NonStop performance metrics are in flight from around the world. These metrics are traveling to multiple computing zones that host Remote Analyst servers. Simultaneously, much larger volumes of data are under analyses by Remote Analyst to look for the unexpected, whether it has already taken place, or is potentially about to happen in the weeks ahead. As you read this, HPE NonStop customers employing Remote Analyst are receiving analyses, alerts, and instructions.

At its base, Remote Analyst enjoys a rather simple architecture. A module is installed on your HPE NonStop servers that collects and forwards performance metrics (no customer data) via secured connections to Remote Analyst servers in the cloud. These metrics are immediately inserted into massive databases, which are run through a battery of analyses. The scope of these analyses can be to simply verify whether the current metrics fit the expected pattern (for similar minutes of day/weekday) when compared to past activity. Other analyses done simultaneously, can take it to a much deeper level.

Every I/O and transaction is analyzed to look for even the smallest data. Exceptions are noted and trend lines are built. Analyses/reports are prepared and emailed to the distribution list as requested or desired.

All of the above happens 24x7, going through petabytes of data, performing millions of instructions/transactions every second, without a single human having to take part. It is a technological feat, but one which also delivers solid value to its customer base.



The internal complex architecture of Remote Analyst also allows for its perceived simplicity. Before joining other metrics and becoming part of Remote Analyst's massive database, each metric goes through at least three servers, two firewalls, and multiple verifications. Remote Analyst uses smart internal algorithms that deliver simple solutions to users. These are:

Resource utilization levels

For every HPE NonStop resource, Remote Analyst questions whether its utilization falls above or below acceptable levels for the type of resource. If it's above threshold levels, we need to question why, who is responsible, and what are the risks? Remote Analyst keeps a database of accepted industry standards and what may be acceptable at each installation.

Companies use their HPE NonStop servers based on their business operational requirements and schedules. This may at times put momentary pressure on certain resources (perhaps at off-hours), which may not meet the accepted industry levels. However, Remote Analyst has the intelligence to know if this fits a pattern that has proven not to interrupt business for this specific server and application. For example, batch jobs that need to run at high resource consumptions, while business applications are running at low utilization or are offline.

Load balancing analyses

Another question is whether resource utilization is balanced across available resources. If there are any imbalances—is that acceptable, or are there application limitations that require it? Maybe there is a single batch process that runs hot and the excessive load is due to application design?

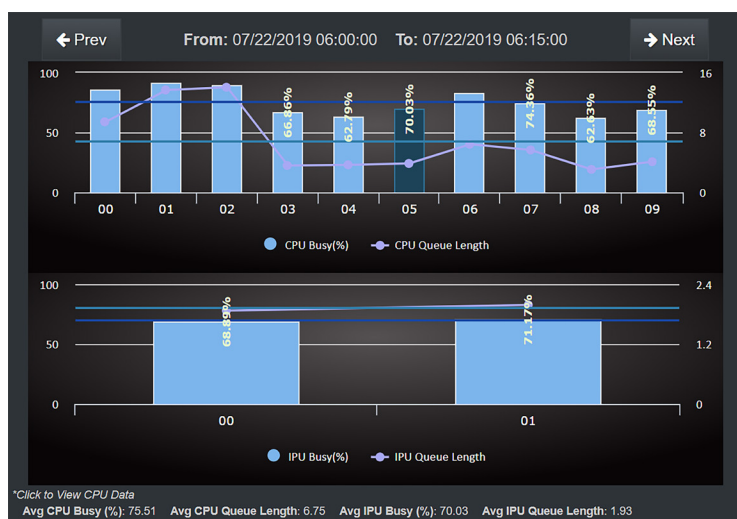


Figure 1. CPU data

The analysis of load balancing is not just about a higher-level view where all applications are put in the same basket and viewed as one. This could be a valid analysis method if all applications and processes are of the same importance, but they rarely are the same. First, there needs to be a priority setting at the application level.

Should application A take priority over application B when they compete for the same resources? Furthermore, within an application, are there certain key modules that need to handle transactions as fast as possible and as soon as they come in? Fortunately, all HPE NonStop servers come with an integrated priority setting design and load-balancing features are incorporated into their architecture.

The next question is whether system resources are available and performing at the desired level. Again, this points to the real-world situations that affect customers. Here is where Remote Analyst demonstrates its value by highlighting specific processes, which may be at or close to reaching the thresholds. Deep analyses are performed to follow each transaction through its internal logic as it seeks system resources, and whether it has to sit in queues.

Note

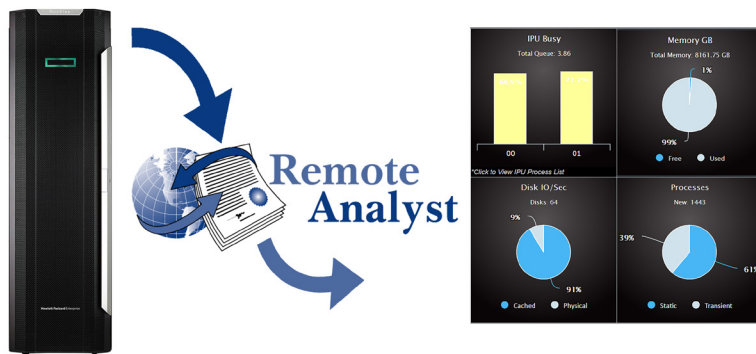
A single transaction may end up sitting in multiple queues for multiple resources (for example, multiple files). Whereas queue rates on each resource may be applicable, the level of queue and wait times experienced by a transaction may not be applicable.



Recommendations and expected results

What should one do to improve performance? What are the exact steps/commands to implement them, and what benefits can be achieved? This is an area where expert analyses and deep knowledge of HPE NonStop servers is very beneficial. Experts have learned over the years that it is very simple to amplify an existing issue by implementing the wrong solution. To determine the best fix, one needs to replay the set of transactions as if the fix was already implemented and then see the result or effect. Furthermore, different scenarios can be considered that include the multiple possible solutions that may apply.

Looking deeper, we can consider possible improvements in other areas. Each solution needs to be examined on its own merits, we also need to examine how a solution will behave when any (or all) of the proposed changes are implemented. The list of possibilities and analytics simply make it impossible for human effort to perform such analyses. It can easily take weeks or months and be completely impractical in today's business environment. Remote Analyst, through use of its high bandwidth and elastic network of computers, performs such analytics every single day. Remote Analyst offers fine-tuned analyses, possible results, specific instructions, and the system-level commands to implement them.



Alerts and links to details

Remote Analyst highlights areas requiring attention with multi-level links to their root causes. Knowing that something is not right is only the beginning. Determining the root cause is key. This is where the nearly unlimited drill-down features of Remote Analyst are very beneficial.

Going from top-level factors such as high CPU busy, an analyst can drill down to processes running on that CPU, how much memory each process consumed, its history, opened files, I/O rate, CPU busy, and memory consumption throughout its life history. Within Remote Analyst, opening new paths is possible at every level. For instance, from a list of files opened by a process, one may open a new path by clicking on a file to see which processes were open, and the file's history over time. At any point in time, you may open yet another path by clicking on any of the processes associated with that file to understand its activities and details.

Remote Analyst does this work via analyses that it performs on its own and it also provides you with an easy to use graphical interface. No typed commands are necessary, but rather a graphical interface to scan and seek deep into lower layers of a system. This aids your analysis in three important areas:

- Enabling human experts to look for root causes on their own.
- Saving significant time via the smart and flexible UI.
- Requiring no command level or specific HPE NonStop server knowledge.

Better or worse than before

For a similar time frame, Remote Analyst answers the question "are we doing better or worse than before?" Remote Analyst also answers "what accounts for the difference?" and "what should be done about it, if anything?" Remote Analyst uses machine learning methods to learn the behavior of your applications and the computing resources they require. It also measures the growth of your business activities.

Multiple years of trend analyses

On average, meaningful long-term trend analyses require a 2-year review. Remote Analyst can trend nearly any metric for any time frame for all available data. This is useful for determining a growth trend and any seasonal changes that exist. For instance, how did the last Christmas season activity differ from the prior year? Given the growth since last year, what type of increased activity should we expect for the upcoming Christmas season? Different businesses may be affected in unique ways during various seasons and holiday periods. Because of these variations, analyses can be performed for separate time frames.

All of this information advances the understanding of system behavior and its applications. It helps to know whether or not they can be relied upon to carry your business during heavy loads without adverse effects on the end-users' experience.



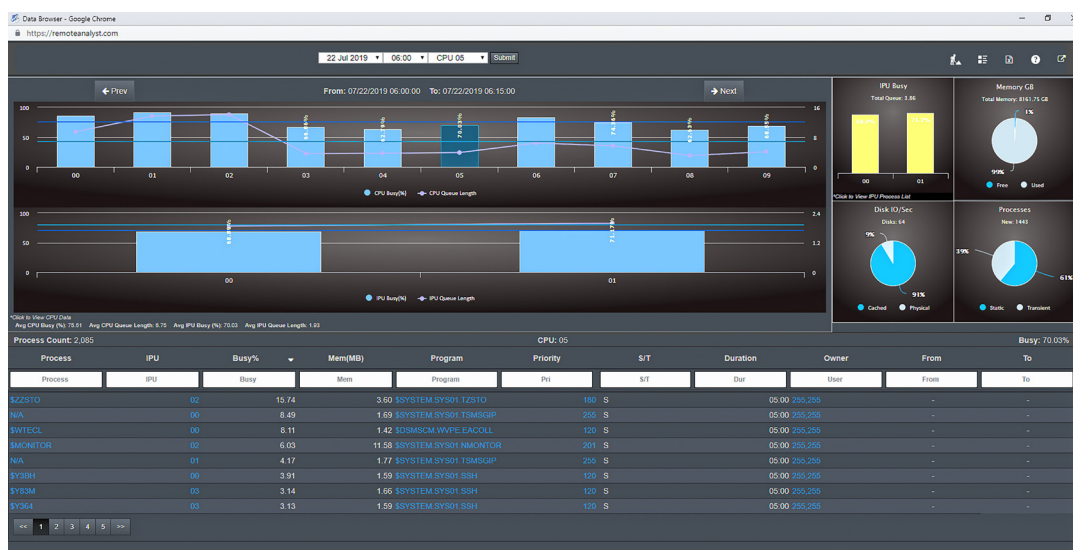


Figure 2. View system details

Summary to detail drill-down options

Remote Analyst allows for multiple to infinite levels of drill down to follow through a chain of related metrics. This feature plays an integral part when experts want to examine the system on their own.

Using the Data Browser feature within Remote Analyst, you can reduce the research time to a fraction of what it may require when using Measure. Here are some key points of what you can do with Data Browser:

- Data Browser is a smart graphical interface that ties in all related system Entities and Metrics. Users can just click through metrics for any of the supported Entities. They can easily find all other related metrics and easily travel through time as if they were doing a system playback.
- Gain advance knowledge of when more resources may be needed.
- Anticipate and correct problems before they affect your business.
- Fine-tune your system to take advantage of all available resource.
- Perform what-if analyses to cover possible anticipated scenarios.

Deep analyses for all resources

(Includes CPU/IPU, disk, process, network, and application)

Remote Analyst looks at anything and everything, including hundreds of detailed analyses and billions of trend analyses. This is where Remote Analyst shines. A series of specific and deep analyses have been developed over years of effort, with the help of HPE NonStop experts from around the world.

In addition, knowledge base and best practices by HPE NonStop performance specialists are reviewed. The result is a series of hundreds of detailed analyses, each of which focuses on a specific area of the server, yet with a link to other related entities.

Remote Analyst users can easily logon to the Remote Analyst site, ask for these analyses on-demand, and receive it via email only minutes later. This is the type of analyses that previously required weeks to prepare via human effort. These analyses are packaged as several workbooks, which includes help and documentation for further support and definition.

To further simplify the process, users can also ask to receive these analyses regularly via email. They simply appear in their inbox, along with other requested analyses such as trends, alert notifications, storage analyses, and more.

What-if analyses

What if the business activity goes up or down? Or what will happen when new and faster processors are employed? What if a new version of the OS is installed? What will happen to our business and what will our utilization look like? Are we ready for it? There are no standard answers to these questions because:

- No two businesses are alike. What works for one company may not be suited for another.
- No two applications are alike. There are often two or more applications (or modules) running on the same server. They can be affected differently as load increases.
- Increase in the business activity does not necessarily translate to increased transactions for all applications. Application loads may be affected differently and some may not be impacted at all.
- More activity can impact a variety of resources differently. Increased transactions may affect storage, network, and CPU/memory, but all at different levels. A new transaction may mean one more data packet coming through the network but may result in (seven) different I/Os on (three) separate drives. Each of these need to be studied separately.
- Different computing resources are not affected in the same way. For similar reasons stated here, each additional transaction may inflict a heavier load on one resource versus another. For example (using arbitrary numbers), CPU usage may increase by 1%, whereas network traffic could increase by 3%.
- Different computing resources have varying threshold levels and are at different utilization levels. For example, a CPU may be at 79% busy and an additional 1% may take it to its first level threshold of 80%. On the other hand, network utilization may be at 40% and a 3% increase will still keep it well below its threshold of 60%.

The main point is that there are several factors to be considered when doing what-if analyses to prepare for an event or a seasonal shift. These factors play out differently at different installations and require access to detail and historical metrics to perform complete analyses. This is exactly what is available in Remote Analyst.

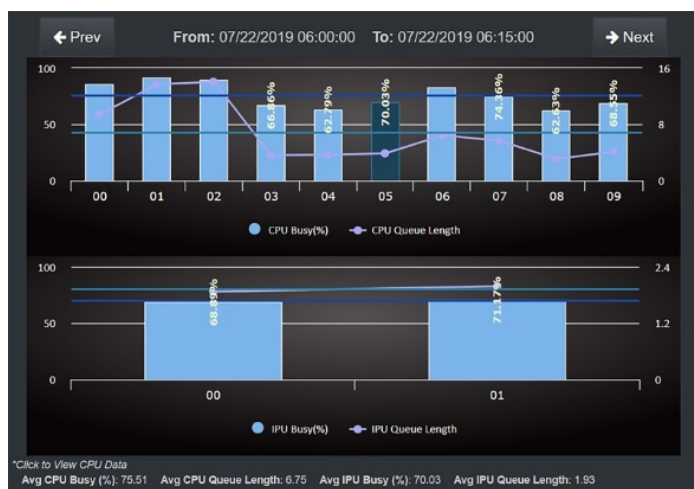


Figure 3. CPU/IPU busy and queue length

Memory leak detection

We just had a new OS or Application release. Is there any unexpected memory leakage, which may affect us (days or weeks) down the road? Memory leaks are another one of those underlying issues, which can easily go undetected until it is too late. A process starting at a 50 MB memory allocation could easily request and use substantially more memory, while failing to let go of the resources that it was previously allocated.

This problem can be associated with a new release of a program, where a developer may have made some small modification to the program's internal architecture and how it utilizes the memory. Issues may not be detected during a QA process where the emphasis is placed on a program's functions. Even when stress tests are performed, they are limited to an increased level of load and transactions, not how they perform over weeks and months of continuous use.



Net increase in memory usage in a problematic architecture is usually small and can take a substantial amount of time before they reach significant levels. What makes it substantially more damaging is that object code is used to start multiple copies of processes in a different CPU (to balance CPU load consumption). The unintended but substantially damaging effect of this is that each process is separately increasing its memory usage and each in a different CPU. Given proper load balancing of processes, they are each building memory usage at about the same level, all hitting thresholds at about the same time and on different CPUs. Potentially this means that multiple CPUs and processes are affected, along with every other process running on those CPUs. All along, CPU busy would have shown to be within acceptable levels and memory usage was on a gradual increase, which makes it even harder to detect any specific object code before it is too late.

Remote Analyst can easily detect such scenarios. It gives users the option of specifying any length of time for analyses, filter out unnecessary scopes (unaffected programs), and look for the culprits that affect the system most. Users can specify minimum memory allocations, static or dynamic. All this can be done via graphical means based upon actual historical data up to the present time.

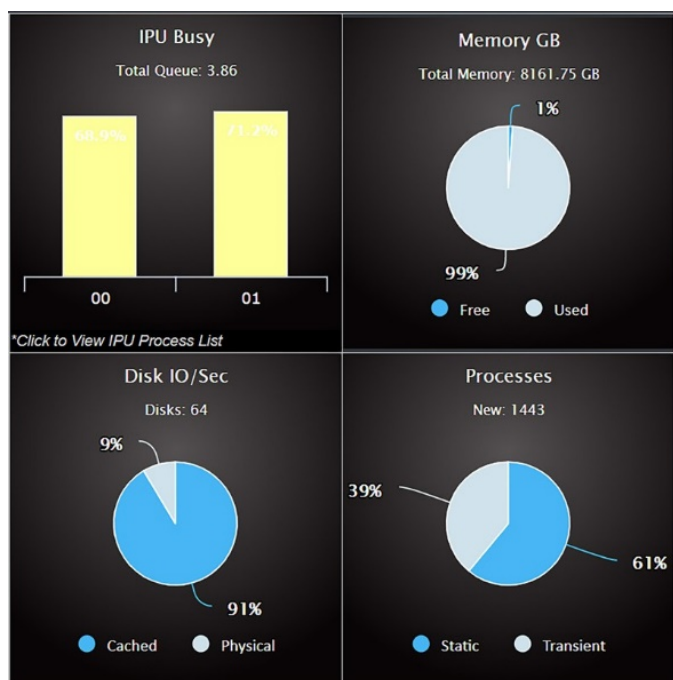


Figure 4. IPU, Memory, Disk I/O, and Processes Charts

Storage capacity analyses and predictions

Do I have enough capacity now? Will I have enough 90 days from now? Are there any disks near capacity now or in the near future?

More transactions pass through CPUs than ever before and nearly every one of them end up in a disk drive as new records. Whether it is a new transaction that needs to be recorded, logs which need to be created, or new bills to keep track of, all require more and more storage.

During technology presentations, we often hear how nearly all of the recorded information in human history came about in the past few years and continues going forward at an accelerated rate. More records, files, and disks need to be monitored and managed, a lot more. Again, this is an area where human effort alone will not be able to keep up.

Remote Analyst scans all HPE NonStop disk drives every night to record storage-level activities (growth, file counts, and more). This information is then sent from the HPE NonStop server to Remote Analyst servers, which is compared against the past activity. This forms the basis for storage growth and capacity analyses. Not only does Remote Analyst show current usage, it can predict (based on past growth) when each volume is set to cross certain thresholds but also shows usage by sub-volume, group, and user IDs.



Detail network analyses

This analysis follows a data byte from its origination to its destination. It also detects and records, which gateways were traveled through from end to end. Analyses reveal whether there was enough bandwidth and what other processes competed with the data en route. Remote Analyst monitors and records all network traffic via Expand and TCP/IP as they go through the HPE NonStop server. Thresholds can be set with visual notifications. Users may also playback any recorded activity, stop at any point in time, and seek more details.

It is important to note, performance analyses are no longer about simple CPU, memory, queue, or any single resource monitoring. Much deeper analyses need to be done, much faster, with built-in behavior analyses, and all offline from the HPE NonStop server.

A by-product and valuable feature of the new architecture offered by Remote Analyst is that users do not have to install any client application on their device. All answers are directed to them via email and when they access the cloud directly via [remoteanalyst](#).

Our Remote Analyst customers are always certain to use the latest available set of software and enhancements. All analytical and UI releases are cloud-based, therefore, available to its worldwide community the moment they are released with no additional requirement placed on users.

The solution is available now

We started by saying that faster computers have allowed us to package more activity over a smaller number of resources, time, and space. This creates a concern where a simple issue, unnoticed, under the surface, can become a major issue over time, and at an accelerated pace. This is a reality at every data center, small or large, that gets progressively challenging as computers and networks become faster.

Deep and consistent performance analyses has become a necessity that is here to stay. This is the type of analyses that cannot be performed by one person, or even an army of human effort. It is impossible to view from any one or more online monitor screens and is best performed offline from the HPE NonStop server.

We have promoted the idea that such a solution was impossible to implement until recently due to unavailability of required technology. Once all the requirements were available, further significant work was needed to design, develop, and test all components to be part of an elastic cloud-based analytic architecture for HPE NonStop servers. Time was also needed for a new solution to prove itself in the real world, and to go through a process of maturity.

Remote Analyst is the answer for performance analyses in the 21st century. It is a mature, proven solution and it is available today.

Here are the key benefits that Remote Analyst offers:

Early warning/outage avoidance

Problems are often detected and reported long before they would surface as an alarm on a monitor screen. This enables system managers to study the issues, craft, and implement a solution (on their own or one recommended by Remote Analyst) before business activities and customers are affected.

Reduced cost

Low-cost license fees for Remote Analyst and several other key cost factors:

- Preventing an unscheduled outage and its consequences (immediate and future) could be the greatest cost-saving benefit of using Remote Analyst. Early warning methods can eliminate or significantly reduce downtimes attributed to improper resource usage and/or inadequate configuration/design. Problems are detected and reported long before they can affect the business.
- Low-cost cloud-based computing resources (at no additional cost) are used for all analytics. This removes the requirement (and the cost it bears) for a customer to allocate HPE NonStop server resources (computing and storage) for deep performance analyses.
- Full automation means near-zero human effort required to do the analyses. This too reduces cost and increases productivity.
- Less human effort needed means less training requirements. This further reduces cost.

Less management requirements

Remote Analyst operates proactively. It informs users if everything is working properly or when problems are on the horizon. Your staff does not need to schedule a time to manually evaluate system performance regularly. It is done for them every hour of every day.



Use our knowledge to your advantage


The HPE NonStop performance experts have been using Remote Analyst for their own Professional Services business solutions for years. They have in-depth knowledge of the Remote Analyst analyses and reporting capabilities. Remote Analyst offers its users the ability to optionally choose whether experts in HPE NonStop consulting organization should receive copies of your systems' performance analyses.

This partnership with HPE NonStop consultants is also an inherent part of the Remote Analyst architecture. It enables HPE performance specialists, upon your choice, to help monitor your NonStop resources. They can immediately notify you of impending or present issues, and engage with you to proactively address those issues.

We look forward to helping you discover the very best from your HPE Integrity NonStop server solutions.

Learn more at

hpe.com/services/consulting



 **Share now**

 **Get updates**



© Copyright 2019 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

iPhone is a trademark of Apple Computer, Inc. registered in the U.S. and other countries. All other third-party marks are property of their respective owners.

a50000053ENW, August 2019