

# Assessing Hardware Failures on IBM PureData System for Analytics N2001, N2002, and N3001

## Overview

This IBM Analytics Support Wed Doc describes possible hardware component failures on IBM PureData System for Analytics, their effect on the system, and its criticality and consequences if multiple components fail.

This document also describes the architecture of IBM PureData System for Analytics to clarify the relationships between components. It also provides guidance on replacement scenarios. This information is intended to help users correctly respond to failures and open problem management records (PMR) with an accurate level of severity.

This document provides information to help you better understand what is needed to perform a part replacement. By using this information, you can plan ahead without needing more information about requirements from an IBM engineer. This document also helps you to identify the effect of a failure on the system and schedule a replacement.

## Notes

This document does not apply to the following IBM PureData System for Analytics models:

- IBM Netezza 100 (Skimmer®) models
- IBM PureData System for Analytics N3001-001

These models feature a different architecture than the architecture that is described in this document.

There is a slight difference between the IBM PureData System for Analytics N1001 models and IBM PureData System for Analytics N2001 and N3001 models.

The IBM PureData System for Analytics N1001 (also known as IBM Netezza TwinFin® ) includes SAS switches, which are not covered in this document. SAS switches were replaced in later models by a cable that directly connects the SPU to the disk enclosure, which is the only difference between the systems.

Figure 1 shows the components and how critical they are to the system. Each component is described in this document.

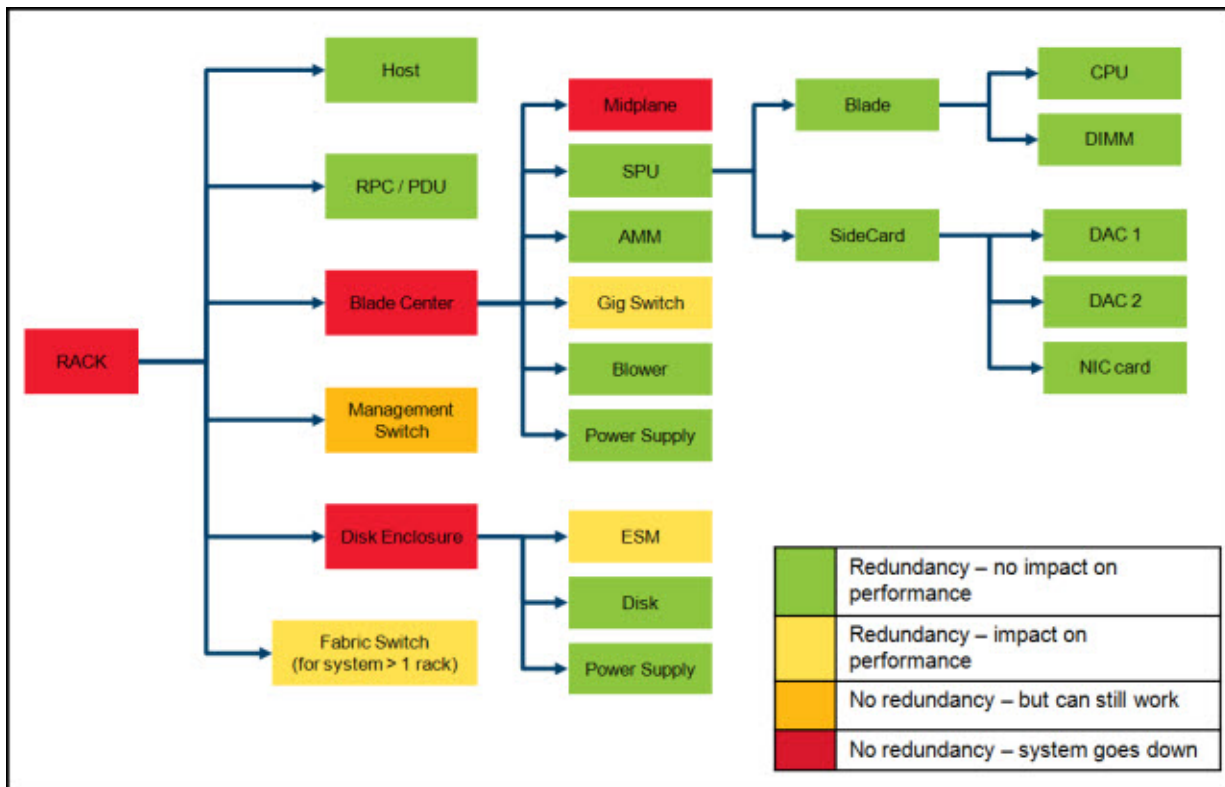


Figure 1. Hardware components and the effect of their failure to the system

The following acronyms are used in Figure 1:

- Remote Power Controller/Power Distribution Unit (RPC/PDU)
- Snippet Processing Unit (SPU)
- Advanced Management Module (AMM)
- Environmental Services Module (ESM)
- Central Processing Unit (CPU)
- Dual in-line Memory Module (DIMM)
- Database Accelerator Card (DAC)
- Network Interface Controller (NIC)

In an N200x / N3001 system, an SPU runs at 100% performance when it handles no more than 40 dataslices. Because each system has an "extra" SPU, each SPU handles fewer than 40 dataslices.

In an N200x/N3001-002 system, 40 dataslices are available and only two SPUs. Therefore, each SPU handles only 20 dataslices. When an SPU fails, the other SPU must manage 40 dataslices and the difference is more significant.

This document describes the following failures:

- Disk
- SPU
- ESM
- Disk enclosure power supply
- BlaceCenter power supply
- BladeCenter blower
- BladeCenter gig switch
- Management switch
- Host disk
- Host
- Rack power supply

### **Disk failure**

This type of failure is the most commonly seen failure on your system. How often this failure occurs depends on the use of the database. It is normal to see failures each month or even each week on larger systems that are heavily used. Failure and replacement are done while the database remains online.

After a disk failure, a regeneration process is started that copies data from the remaining dataslice copy to a spare disk. Performance is decreased to 40% of original value until the regeneration process completes.

The regeneration process is not the only process that slows down the system as it takes only 10% of overall performance. The other 50% of performance degradation is the result of the difference in the number of disks that are available before and after a failure. Before the failure, two disks deliver data for each dataslice. After a failure, a single disk delivers data for two dataslices, which causes the 50% performance degradation.

It is this issue of one disk being available that is the drawback of parallel systems. Overall performance is only as fast as the slowest component.

The time a regeneration process runs depends on the amount of data it must copy.

The copy time is 60 megabytes per second (MBps) on an idle system and 10 MBps on a system under use. For example, if dataslices hold on average 100 GB of data, it takes approximately 10.000 seconds to copy one dataslice on a system that is used (regeneration speed of 10 MBps). Because two dataslices are on each disk, both dataslices must be copied. The total time that is needed to complete this copy is approximately 20.000 seconds (6 hours).

Disk failures are a concern only if spare disks are unavailable.

### **SPU failure**

SPU failures occur occasionally and are not seen often.

Several components on a SPU can fail and regardless of which component fails, it is the SPU that fails. (For this document, it does not matter which component causes the SPU failure.)

The database is paused when a SPU fails, which interrupts any insert queries and select queries. The pause lasts for approximately 20 minutes, but can be longer if the system does not detect any particular error and attempts to power cycle the SPU to restore it to normal operation.

Although there are situations in which a SPU fails and appears to be unresponsive, a power cycle often clears the issue and there is no need for replacements.

When the SPU fails, replacements can be made without any outage on the system. Because the SPU is not used, it can be pulled from the system without affecting any other part of the system.

After the replacement is complete, it takes an outage of approximately 15 minutes to return the SPU to active operation. This outage can be done at any time by running the `nzhw rebalance` command as `user nz`. It does not require onsite assistance.

### **ESM failure**

The ESM is a communication and monitoring component. Each disk enclosure includes two ESMs. Its failure does not disable the database because communication is possible through the second ESM. Performance might be affected, depending on which ESM fails.

If performance is degraded, it often is not more than 50%.

Replacing the ESM is an online process and does not require that the database is disabled. No rebalance must be done and after the new ESM is operational, new paths are discovered automatically and used by the system.

### **Disk enclosure power supply failure**

A disk enclosure has two power supplies and can be fully functional on a single supply.

The failure does not cause any outages and part replacement is an online process.

Although performance is not affected, the power supply must be replaced because a failure of the second power supply disables the entire database and no processing can be done until the disk enclosure becomes operational.

### **IBM BladeCenter power supply failure**

There are four power supplies in a BladeCenter. The power supplies are divided into two power domains.

If two power supplies from the same power domain fail, half of the blades cannot power and the system might fail (or at least run at 50% performance).

A BladeCenter's power supply failure does not cause any outage and its replacement is an online process.

### **IBM BladeCenter blower failure**

There are two blowers that cool the blades in the BladeCenter. If one blower fails, the second blower automatically runs at full power, which is easily heard in the datacenter.

Although the single blower provides enough cooling to sustain system usage, it is important to replace the failed blower quickly. SPUs can overheat, which can result in the database and all of the elements of the SPUs failing.

### **IBM BladeCenter gig switch failure**

The gig switch provides connectivity between the BladeCenter and the hosts. It is the main data transfer backbone between SPUs. Because there are two switches in each BladeCenter, you still have connectivity and the system runs with only one switch working.

The decrease in query execution varies because it does not affect the CPUs and their processing power, only data transfer. In a worst-case scenario, the effect is a decrease in performance of 50%. In reality, you see a less dramatic decrease of performance. For small systems, this decrease might be 0%.

When the switch fails, it might trigger a restart of the database. Replacement requires an outage of approximately three hours.

### **Management switch failure**

This failure is tricky because the system remains online despite the host being unable to communicate with various components, such as RPCs and AMMs. Therefore, the system cannot power cycle SPUs, perform a host failover, or receive confirmations that SPUs completed starting. If the database pauses or stops, it cannot be brought back up until the failed management switch is replaced.

It is critical to replace the switch as soon as possible. Failures of this component are rare, but must be treated with critical attention when they occur.

The replacement requires a four-hour system outage.

### **Host disk failure**

The hosts use RAID 10 and include a spare disk. Depending on the model, seven disks are used (three mirror pairs of two and a spare) or five disks (two mirror pairs and a spare). Failure of a single disk is not cause for alarm because the spare disk takes over and RAID 10 is preserved with data mirroring.

Multiple host disk failures can be serious problem. If the system loses both copies of the data, the host fails and must be rebuilt. The rebuilding process requires a 10-hour system outage. Replacing a host disk is an online process.

### **Host failure**

A total host failure can be the result of many issues, such as a motherboard or RAID controller failure.

Whatever the cause, your system is affected only if the issue occurs on the active host. Because the standby host is on standby and database connections are not running through the host, failure does not affect query processing.

If the active host fails, the database is moved from the failed host to the standby host so that processing can resume. However, this move interrupts all processing and database connections. After 15 minutes, the database often is back online and running on the second host, ready to resume running queries.

The outage that is needed to fix a host depends on the components that must be replaced. There is no performance effect on the environment; however, you cannot have a high availability system until the host is repaired. If the last remaining host is lost, all of the data that is associated with that host also is lost.

### **Rack power supply failure**

When we describe a rack power supply failure, we are referring to each rack. Each rack features its own independent power from other racks. The one exception is N3001/N200x-002 models, whose racks are referred to as a quarter rack. These racks require less power and are powered by two long RPCs that are directly connected to the datacenter's power. One RPC represents the rack's left side power and the other the right-side power.

The following types of power supplies provide power to the entire rack:

- From your datacenter, there are two power cables that are connected to the rack. These cables are from different power domains to sustain redundancy and are connected to the PDUs.

The PDUs are unmanageable and the system does not monitor them. Although it is rare for a PDU to fail, a failure is easy to identify because an entire side of a rack does not have power.

- The second type of power supply is the RPCs, which are monitored. By logging in to the RPCs, any socket on that power supply can be and powered down or up.

The power layout features two power cables from the datacenter that are connected to two PDUs. Each PDU supplies power to two RPCs and the BladeCenters.

Power supplies fail rarely. A failure often is the result of a problem with a breaker or a loose power cable.

Each RPC supplies power to the disk enclosures, management switch, and hosts. Replacing an RPC or PDU requires a two-hour system outage.

### Failure cheat sheet

The types of failures and their effect on the system are listed in Table 1.

Table 1 Failure cheat sheet

Failure type	Effect	Result
Disk	Failure causes outage	No
	Replacement needs outage	No
	Impact on performance	60% lower until regeneration finishes
	Criticality	Minor
SPU	Failure causes outage	Yes
	Replacement needs outage	Only 15 minutes after replacement is completed
	Impact on performance	None (unless an IBM PureData System for Analytics N2001-002, N2002-002, or N3001-002 system is used)
	Criticality	Minor
ESM	Failure causes outage	No
	Replacement needs outage	No
	Impact on performance	Up to 50%
	Criticality	Moderate
Disk enclosure power supply	Failure causes outage	No
	Replacement needs outage	No
	Impact on performance	None
	Criticality	Minor
IBM BladeCenter power supply	Failure causes outage	No
	Replacement needs outage	No
	Impact on performance	None
	Criticality	Minor
IBM BladeCenter blower	Failure causes outage	No
	Replacement needs outage	No
	Impact on performance	None
	Criticality	Moderate
IBM BladeCenter gig switch	Failure causes outage	Yes (but not always)
	Replacement needs outage	Yes (three hours)
	Impact on performance	Up to 50%
	Criticality	Moderate
Management switch	Failure causes outage	No (but system cannot start)
	Replacement needs outage	Yes
	Impact on performance	None
	Criticality	Critical

Host disk	Failure causes outage	No
	Replacement needs outage	No
	Impact on performance	None
	Criticality	Minor
Host	Failure causes outage	Only if active host fails
	Replacement needs outage	No
	Impact on performance	None
	Criticality	Moderate
Rack power supply	Failure causes outage	No
	Replacement needs outage	Yes (two hours)
	Impact on performance	None
	Criticality	Moderate

### More information

For more information, see the following resources:

- IBM Knowledge Center:  
[https://www.ibm.com/support/knowledgecenter/SSULQD/SSULQD\\_welcome.html](https://www.ibm.com/support/knowledgecenter/SSULQD/SSULQD_welcome.html)
- IBM PureData System:  
<https://www.ibm.com/software/data/puredata/analytics/>

### Special notices

This document does not apply to the Netezza-100 (also known as Skimmer) and the N3001-001.

PureData for Analytics N1001 systems are slightly different because they use SAS switches, which are not covered in this document.

# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

**© Copyright International Business Machines Corporation 2016. All rights reserved.**



This document was created or updated on September 29, 2016.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:  
[ibm.com/redbooks](http://ibm.com/redbooks)
- Send your comments in an e-mail to:  
[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)
- Mail your comments to:  
IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips1345.html> .

## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

- IBM®
- IBM PureData®
- IBM z™
- PureData®
- RACF®
- Redbooks®
- Redbooks (logo)®
- System z®
- z Systems™
- z/OS®

The following terms are trademarks of other companies:

Netezza, Skimmer, TwinFin, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.