

PerfCap Corporation
Performance and Capacity Planning Services

**XDEV Shared Infrastructure Dev/Test
Capacity Planning Report
February 17, 2011**

Prepared by:
R. Grant
PerfCap Corporation



Preface

This document presents the findings of PerfCap Corporation's capacity monitoring of XDEVs' Shared Infrastructure Dev/Test servers from December 2010, through mid-February, 2011. The purpose of this report is present significant findings regarding resource capacity utilization projections 60 days into the future.

Disclaimer

The report and recommendations are based on information collected on XDEV's Shared SQL (Dev/Test), Shared Oracle (Dev/Test), Shared Oracle (OAT), and Shared Solaris (Dev/Test) servers. XDEV is responsible for evaluating PerfCap's recommendations. The services furnished are in the nature of professional assistance and advice. PerfCap warrants that the services will be performed in a workman-like manner. However, PerfCap does not warrant or guarantee in any form or manner the results and/or achievements of the service provided.

There is no commitment to provide a system, which performs according to the results of this service. The service makes no predictions about capacity requirements for future, unobserved workloads. PerfCap is not responsible for the work or activity of any non-PerfCap personnel. Except for the express warranties stated herein, PerfCap disclaims all warranties, including all implied warranties of merchantability and fitness for a particular purpose, and the stated express warranties are in lieu of all obligations or liabilities on the part of PerfCap arising out of or in connection with the performance of the service.

Notwithstanding the terms of PerfCap's standard terms and conditions, PerfCap's total liability to the customer for any cause whatsoever is limited to the amount paid to PerfCap for the services furnished under this agreement regardless of the form of action.

Table of Contents

PREFACE	I
DISCLAIMER	I
TABLE OF CONTENTS	II
EXECUTIVE SUMMARY	3
CAPACITY FINDINGS.....	3
RECOMMENDATIONS	3
RESOURCE UTILIZATION RISK ANALYSIS	4
CPU UTILIZATION	5
<i>Overview</i>	5
<i>Shared Oracle (Dev/Test) REUXEUUS471, REUXEUUS472</i>	5
<i>Shared Oracle (OAT) REUXEUUS504</i>	6
<i>ESX Host, REUXEUVM181</i>	8
MEMORY UTILIZATION	9
<i>Overview</i>	9
<i>Shared SQL (Dev/Test) - XDEV1RSDIS315, XDEV1RSDIS316, XDEV1RSDIS317, XDEV1RSDIS373</i> ..	10
<i>Shared Oracle (OAT) – REUXEUUS504</i>	11
STORAGE SYSTEM PERFORMANCE	12

Executive Summary

The following summary details the hardware resource utilization trends of the Shared SQL (Dev/Test), Shared Oracle (Dev/Test), Shared Oracle (OAT), and Shared Solaris (Dev/Test) servers.

Capacity Findings

PerfCap PAWZ automated monitoring and trending of hardware resource utilization from November, 2010, through mid-February 2011, provided the following insights:

1. Two servers are at risk of severe CPU congestion. *Shared Oracle (Dev/test) node REUXEUUS471 is at risk of exhausting CPU resources during the first half of March 2011. Shared Oracle (OAT) node REUXEUUS504 is at risk of exhausting CPU resources by the end of February 2011.* The majority of the remaining server CPUs are underutilized.
2. Shared SQL (Dev/Test) node *XDEV1RSDIS373 appears to possess a memory leak. If not checked, this node will exhaust its memory by March 6, 2011.* Historically this node has experienced a gradual memory leaks. These leaks have been “flushed” at various intervals by some unknown event, possibly a system or application failure or reboot. The remaining servers show no prospect of memory starvation at this time.
3. Shared Solaris (Dev/Test) nodes REUXEUUS550, REUXEUUS551, REUXEUUS558 and REUXEUUS560 and Shared Oracle (Dev/Test) nodes REUXEUUS638 and REUXEUUS639 experience consistent poor disk response time.

These development and test servers do not directly impact performance of production applications. Nevertheless, because the same risks may propagate into a production environment at a future date, XDEV should examine and clearly understand each issue.

Recommendations

PerfCap makes the following recommendations:

1. Monitor and balance loads between Shared Oracle (Dev/test) nodes REUXEUUS471 and REUXEUUS472.
2. Monitor and balance loads between Shared Oracle (OAT) nodes REUXEUUS503 and REUXEUUS504.
3. Determine the *root* cause of the memory leak on Shared SQL node XDEV1RSDIS373.
4. Consider upgrading storage provisioning for poorly performing volumes on Shared Solaris (Dev/Test) nodes REUXEUUS550, REUXEUUS551, REUXEUUS558 and REUXEUUS560 and Shared Oracle (Dev/Test) nodes REUXEUUS638 and REUXEUUS639.

Resource Utilization Risk Analysis

Table 1 lists the nodes monitored in the SI Dev/Test group. PerfCap PAWZ software captured data on these servers beginning in early December 2010. This data was archived in the PAWZ database and resource utilizations and performance metrics trended and projected day to day since that time.

<i>Node Name</i>	<i>Purpose</i>	<i>Model</i>	<i>Logical CPUs</i>	<i>Memory</i>
XDEV1RSDIS315	SI - SQL Server 2005 - DEV TEST (#1)	ProLiant DL585 G2	8	64 GB
XDEV1RSDIS316	SI - SQL Server 2005 - DEV TEST (#2)	ProLiant DL585 G2	8	64 GB
XDEV1RSDIS317	SI - SQL Server 2005 - DEV TEST (#3)	ProLiant DL585 G2	8	64 GB
XDEV1RSDIS373	SI - SQL Server 2005 - DEV TEST (#4)	ProLiant DL585 G2	8	64 GB
REUXEUUS471	SI - Oracle 10g - DEV TEST (#1)	Sun Netra-T12	24	192 GB
REUXEUUS472	SI - Oracle 10g - DEV TEST (#2)	Sun Netra-T12	24	192 GB
REUXEUUS638	SI - Oracle 10g - DEV TEST (#3)	Sun T5240	128	64 GB
REUXEUUS639	SI - Oracle 10g - DEV TEST (#4)	Sun T5240	128	64 GB
REUXGBUX058	SI - Oracle 10g (x86) - DEV TEST (#1)	Sun i86pc	16	128 GB
REUXGBUX059	SI - Oracle 10g (x86) - DEV TEST (#2)	Sun i86pc	16	128 GB
REUXGBUX060	SI - Oracle 10g (x86) - DEV TEST (#3)	Sun i86pc	16	128 GB
REUXGBUX061	SI - Oracle 10g (x86) - DEV TEST (#4)	Sun i86pc	16	128 GB
REUXEUUS503	SI - Oracle 10g - OAT (#1)	Sun-Fire-V490	8	64 GB
REUXEUUS504	SI - Oracle 10g - OAT (#2)	Sun-Fire-V490	8	64 GB
REUXEUUX259	SI - Oracle 10g (x86) - OAT (#1)	Sun i86pc	16	128 GB
REUXEUUX264	SI - Oracle 10g (x86) - OAT (#2)	Sun i86pc	16	128 GB
REUXGBUS056	SI - Oracle 10g - OAT DR (#1)	Sun-Fire-V490	8	64 GB
REUXGBUX090	SI - Oracle 10g (x86) - OAT DR (#1)	Sun i86pc	16	128 GB
REUXEUUS497	SI - Solaris 10 - DEV TEST (#1)	Sun Enterprise-T5220	64	64 GB
REUXEUUS550	SI - Solaris 10 - DEV TEST (#2)	Sun T5240	128	64 GB
REUXEUUS551	SI - Solaris 10 - DEV TEST (#3)	Sun T5240	128	64 GB
REUXEUUS558	SI - Solaris 10 - DEV TEST (#4)	Sun T5240	128	64 GB
REUXEUUS560	SI - Solaris 10 - DEV TEST (#5)	Sun T5240	128	64 GB
REUXEUM181	SI - VMware ESX 3.5 - DEV TEST (#1)	ProLiant BL685c G5	16	128 GB

Table 1 - List of SI Dev/Test Nodes

While numerous performance metrics were collected and recorded, CPU utilization, memory utilization, and disk volume response time have been examined to determine risk of exhausting available resources. Projection trends were based on approximately two months of data and were extended 60 days into the future. While most resources are underutilized, several systems are at risk of exhausting CPU or memory resources. Numerous disk volumes exhibited poor response time. The following sections will detail these observations.

CPU Utilization

Overview

Table 2 lists the monthly average and monthly 1-minute peak CPU utilizations for each SI Dev/Test node. The table is broken into five sections: Shared SQL (Dev/Test), Shared Oracle (Dev/Test), Shared Oracle (OAT), Shared Solaris (Dev/Test), and one ESX host. The CPUs of those nodes listed in purple are consistently underutilized.

Node	November	December	January	February	Average/ Maximum
XDEV1RSDIS315	-	16.3/72.8	7.4/64.9	15.5/96.4	13.1/96.4
XDEV1RSDIS316	-	25.2/77.6	27.1/94.9	10.1/91.8	20.8/94.9
XDEV1RSDIS317	-	17.2/62.2	22.5/76.1	16.7/56.7	18.8/76.1
XDEV1RSDIS373	-	12.3/57.7	15.0/62.4	18.0/65.7	15.1/65.7
REUXEUUS471	65.3/100.0	68.5/100.0	65.1/100.0	74.8/100.0	68.4/100.0
REUXEUUS472	73.9/100.0	68.1/100.0	85.1/100.0	-	75.7/100.0
REUXEUUS638	-	0.8/2.9	2.1/11.8	6.9/17.1	3.3/17.1
REUXEUUS639	-	0.4/2.3	0.4/8.9	0.4/1.4	0.4/8.9
REUXGBUX058	-	35.5/93.1	47.8/85.2	43.5/93.1	42.2/93.1
REUXGBUX059	-	43.0/98.4	45.6/93.2	49.8/78.7	46.1/98.4
REUXGBUX060	-	4.9/57.4	6.7/75.9	5.2/32.5	5.6/75.9
REUXGBUX061	-	22.3/49.7	24.8/97.7	26.8/51.6	24.6/97.7
REUXEUUS503	-	44.2/100.0	41.3/100.0	46.9/100.0	44.1/100.0
REUXEUUS504	-	84.9/100.0	77.0/100.0	84.6/100.0	82.2/100.0
REUXEUUX259	-	39.7/76.0	30.9/92.0	32.2/65.5	34.3/92.0
REUXEUUX264	-	1.6/9.3	0.8/16.2	0.6/5.9	1.0/16.2
REUXGBUS056	-	1.6/17.3	1.6/21.3	1.5/16.8	1.6/21.3
REUXGBUX090	-	12.3/38.3	7.1/38.5	6.1/14.7	8.5/38.5
REUXEUUS497	-	19.7/52.2	22.9/49.1	20.2/29.9	20.9/52.2
REUXEUUS550	-	7.3/17.7	8.9/18.5	8.4/22.9	8.2/22.9
REUXEUUS551	-	4.8/18.7	4.4/27.8	4.0/27.0	4.4/27.8
REUXEUUS558	-	6.0/13.1	7.9/13.2	7.2/10.4	7.0/13.2
REUXEUUS560	-	5.6/13.8	8.6/15.2	8.5/16.4	7.5/16.4
REUXEUVM181	-	70.3/92.9	76.0/93.9	59.8/90.9	68.7/93.9
SI Dev/Test Summary	67.6/100.0	26.5/100.0	26.6/100.0	24.3/100.0	36.3/100.0

Table 2 - Average CPU Utilization

Three nodes exhibited high CPU utilization or a trend towards high utilization within two months:

1. Shared Oracle (Dev/test): REUXEUUS471, REUXEUUS472
2. Shared Oracle (OAT): REUXEUUS504

The remaining nodes' CPUs were underutilized and were not projected to experience CPU congestion in the next two months.

Shared Oracle (Dev/Test) REUXEUUS471, REUXEUUS472

Figures 1 and 2 illustrate the historical CPU utilization of Shared Oracle (Dev/Test) nodes, REUXEUUS471 and REUXEUUS472. Node REUXEUUS471 exhibits a long-term trend that suggests substantial CPU congestion will prevail within the first two weeks of March. The most recent data for REUXEUUS472 (late January) indicated that a state of CPU resource congestion had already been reached. Rebalancing of load across these two nodes is advised.

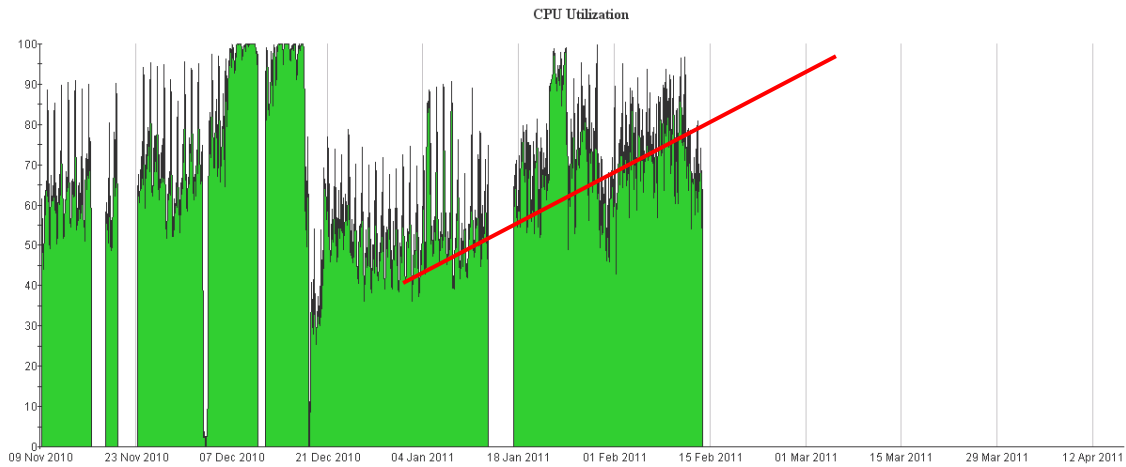


Figure 1 - REUXEUUS471 CPU Utilization Trend

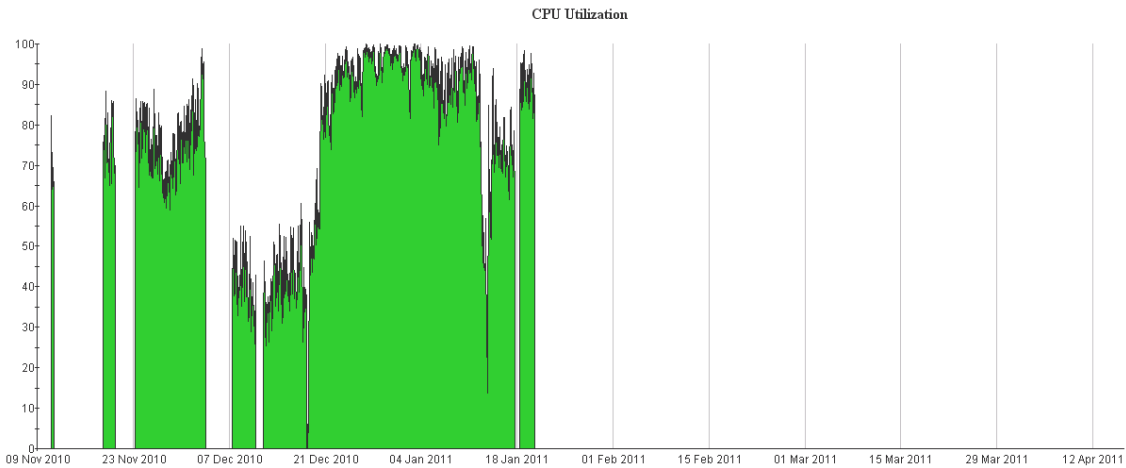


Figure 2 - REUXEUUS472 CPU Utilization Trend

Shared Oracle (OAT) REUXEUUS504

Figure 3 illustrates the trend of CPU utilization for the Shared Oracle (OAT) node, REUXEUUS504. *Since mid-January, this node has experienced a CPU utilization trend, which, if continued, will result in sustained CPU congestion by the end of February.* REUXEUUS504 has eight logical CPUs. Figure 4 indicates the relationship between CPU utilization and CPU queue length on this node. Because there are eight logical CPUs on this node, any time the average queue length exceeds this value, threads experience CPU wait time, or CPU congestion. For example, an average queue length of 48 would mean that the average thread will spend five times what it spends using CPU, waiting for CPU. Note that above 95% CPU utilization this situation degrades rapidly. For this reason, and because of the variability of CPU demand, PerfCap recommends that average CPU utilization be limited to no more than 90% on an 8-CPU system.




Figure 3 – REUXEUUS504 CPU Utilization Trend




Figure 4 - REUXEUUS504, CPU Queue Length vs CPU Utilization

ESX Host, REUXEUV181

The Dev/Test ESX Server, REUXEUV181, experiences erratic CPU resource demand indicative of varied virtual machine workloads as seen in Figure 5. Because of the nature of virtual machine provisioning under ESX, high CPU utilization is not generally a concern, provided each VM is insulated in its resource demands from other VMs with the use of resource caps. The greater concern is what is happening within the VMs. This evaluation requires an examination of the performance metrics reported within the VMs.




Figure 5 - REUXEUV181 CPU Utilization Trend

Memory Utilization

Overview

Table 3 lists the monthly average and monthly 1-minute peak memory utilizations for each SI Dev/Test node. The table is broken into five sections: shared SQL (dev/test), shared Oracle (dev/test), shared Oracle (OAT), shared Solaris (dev/test), and one ESX hosts. The memory of those nodes listed in purple was consistently underutilized. The following nodes exhibited high memory utilization or a trend towards high utilization within two months:

1. Shared SQL (Dev/test): XDEV1RSDIS315, XDEV1RSDIS316, XDEV1RSDIS317, XDEV1RSDIS373
2. Shared Oracle (OAT): REUXEUUS504

The remaining nodes were underutilized and none exhibited a trend of memory utilization which indicate memory starvation in the next two months.

Node	November	December	January	February	Average/ Maximum
XDEV1RSDIS315	-	96.6/99.7	69.8/99.7	89.8/100.0	85.4/100.0
XDEV1RSDIS316	-	96.8/99.7	97.0/99.9	46.3/99.9	80.0/99.9
XDEV1RSDIS317	-	95.9/99.2	95.9/98.7	88.3/99.7	93.4/99.7
XDEV1RSDIS373	-	95.3/99.6	95.3/99.8	91.8/99.5	94.1/99.8
REUXEUUS471	58.3/69.2	58.5/78.8	48.9/59.9	58.6/62.9	56.1/78.8
REUXEUUS472	66.3/71.2	65.7/91.0	74.2/86.8	-	68.7/91.0
REUXEUUS638	-	16.2/18.6	21.9/37.1	33.9/36.1	24.0/37.1
REUXEUUS639	-	5.9/10.3	6.5/16.2	8.5/10.6	7.0/16.2
REUXGBUX058	-	49.7/59.5	54.4/70.4	55.3/63.9	53.1/70.4
REUXGBUX059	-	56.9/66.6	60.7/97.9	62.1/70.4	59.9/97.9
REUXGBUX060	-	38.2/47.7	36.4/98.5	35.1/38.2	36.6/98.5
REUXGBUX061	-	47.2/57.8	51.9/63.8	55.2/64.2	51.4/64.2
REUXEUUS503	-	51.3/60.7	49.5/60.5	50.3/71.5	50.4/71.5
REUXEUUS504	-	91.5/98.4	86.5/97.6	91.7/96.9	89.9/98.4
REUXEUUX259	-	35.5/42.9	36.2/44.8	36.9/45.2	36.2/45.2
REUXEUUX264	-	9.3/9.8	7.4/29.1	6.0/7.7	7.6/29.1
REUXGBUS056	-	7.1/9.0	7.7/9.6	7.9/9.9	7.6/9.9
REUXGBUX090	-	28.4/38.6	27.8/39.5	28.5/39.2	28.3/39.5
REUXEUUS497	-	29.6/38.5	34.1/40.2	38.2/41.6	34.0/41.6
REUXEUUS550	-	41.4/51.9	50.1/62.2	56.3/65.9	49.2/65.9
REUXEUUS551	-	23.3/29.9	27.0/36.1	27.0/34.2	25.8/36.1
REUXEUUS558	-	23.4/31.9	26.6/33.1	29.2/32.8	26.4/33.1
REUXEUUS560	-	22.0/34.9	24.9/30.1	26.7/31.3	24.5/34.9
REUXEUVM181	-	33.5/42.0	35.8/50.5	21.6/46.5	30.3/50.5
SI Dev/Test Summary	-	96.6/99.7	69.8/99.7	89.8/100.0	85.4/100.0

Table 3 – Average Memory Utilization

Shared SQL (Dev/Test) - XDEV1RSDIS315, XDEV1RSDIS316, XDEV1RSDIS317, XDEV1RSDIS373

The Shared SQL (Dev/Test) nodes (XDEV1RSDIS315,316,317,373) all experienced high memory utilization. This may be attributed to SQLserver using as substantial portion of physical memory for its internal cache as it should do when memory is readily available.

Figure 6 illustrates a problematic trend of memory utilization in node XDEV1RSDIA373. This trend, if unchecked, will reach 100% memory utilization by March 6, 2011. This growth in memory utilization may be attributed to an apparent memory leak in the processes running the image, wmiprvse.exe (Figure 7) and day-to-day increase in paged pool (Figure 8). Historically this node has experienced a gradual memory leaks. These leaks have been “flushed” at various intervals by some unknown event, possibly a system or application failure or reboot.



Figure 6 - XDEV1RSDIS373 Memory Utilization Trend

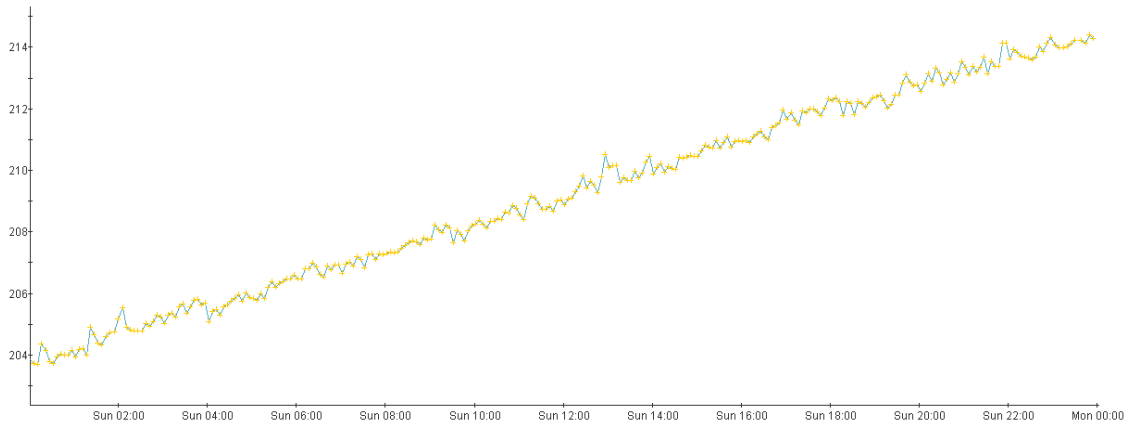


Figure 7 – XDEV1RSDIS373 WMIPRVSE.EXE Per-process Memory Usage (6-Feb-2011)



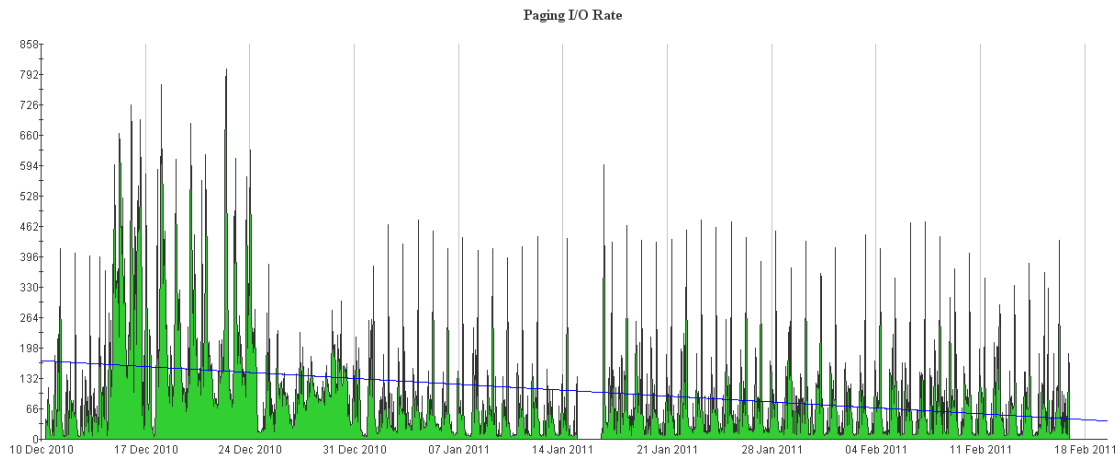


Figure 8 - XDEV1RSDIS373 Paged Pool Growth

Shared Oracle (OAT) – REUXEUUS504

The Shared Oracle (OAT) node, REUXEUUS504 also exhibited high memory utilization. However, there is no indication of an increasing trend of either memory utilization or page fault I/O. This system is using its memory well and not experiencing significant memory resource congestion.





Storage System Performance

The following servers experienced poor disk response times:

1. Shared Solaris (Dev/Test): REUXEUUS550, REUXEUUS551, REUXEUUS558, REUXEUUS560
2. Shared Oracle (Dev/Test): REUXEUUS638, REUXEUUS639

Table 4 lists the servers that consistently experience poor disk response times:

Node	Disks
REUXEUUS550	c1t0d0s0, c1t1d0s0, md10, md11, md12
REUXEUUS551	c1t0d0s0, md10, md11
REUXEUUS558	c1t0d0s0, c1t1d0s0, md10
REUXEUUS560	c1t1d0s0
REUXEUUS638	c1t0d0s0, c1t1d0s0, md10, md11, md12
REUXEUUS639	c1t0d0s0, c1t1d0s0, md10, md11, md12

Table 4 - Nodes, Disks with High Response Time

Response times greater than 20ms are considered poor. These poorly performing disks did not have inordinately high I/O and data rates. Therefore the disks/LUNs and/or their controllers are performing poorly because of inherent limitations or load from other sources.