

Performance Monitoring Process for Out of Standard Applications

Prepared by John Van Wagenen for the
34th Annual International Conference of the
Computer Measurement Group, Inc.
Las Vegas, Nevada USA
Tuesday Dec 9, 2008 2:45 PM




Systems+Processes Division

Caterpillar: Public Blue

CATERPILLAR®
TODAY'S WORK. TOMORROW'S WORLD.™

Does your utilization/capacity picture look like this?



Very Very Busy

We call this a Performance Incident (it lasted for 6 hours)

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

and now a word from my sponsor

Caterpillar is . . . Making Progress Possible

- World's largest manufacturer of construction and mining equipment, diesel and natural gas engines, and industrial gas turbines
- #1 or #2 for each industry we serve with a 2010 goal of market leadership for every major product group on every continent
- Caterpillar Inc. was established in 1925 with the merger of Holt and Best tractor companies and is located in Peoria Illinois, USA.

Global Enterprise

Over 96,000 employees

Operations in 40 countries

Dealer Community

Over 100,000 employees

Spanning 200 countries

Systems+Processes Division

Caterpillar: Public Blue



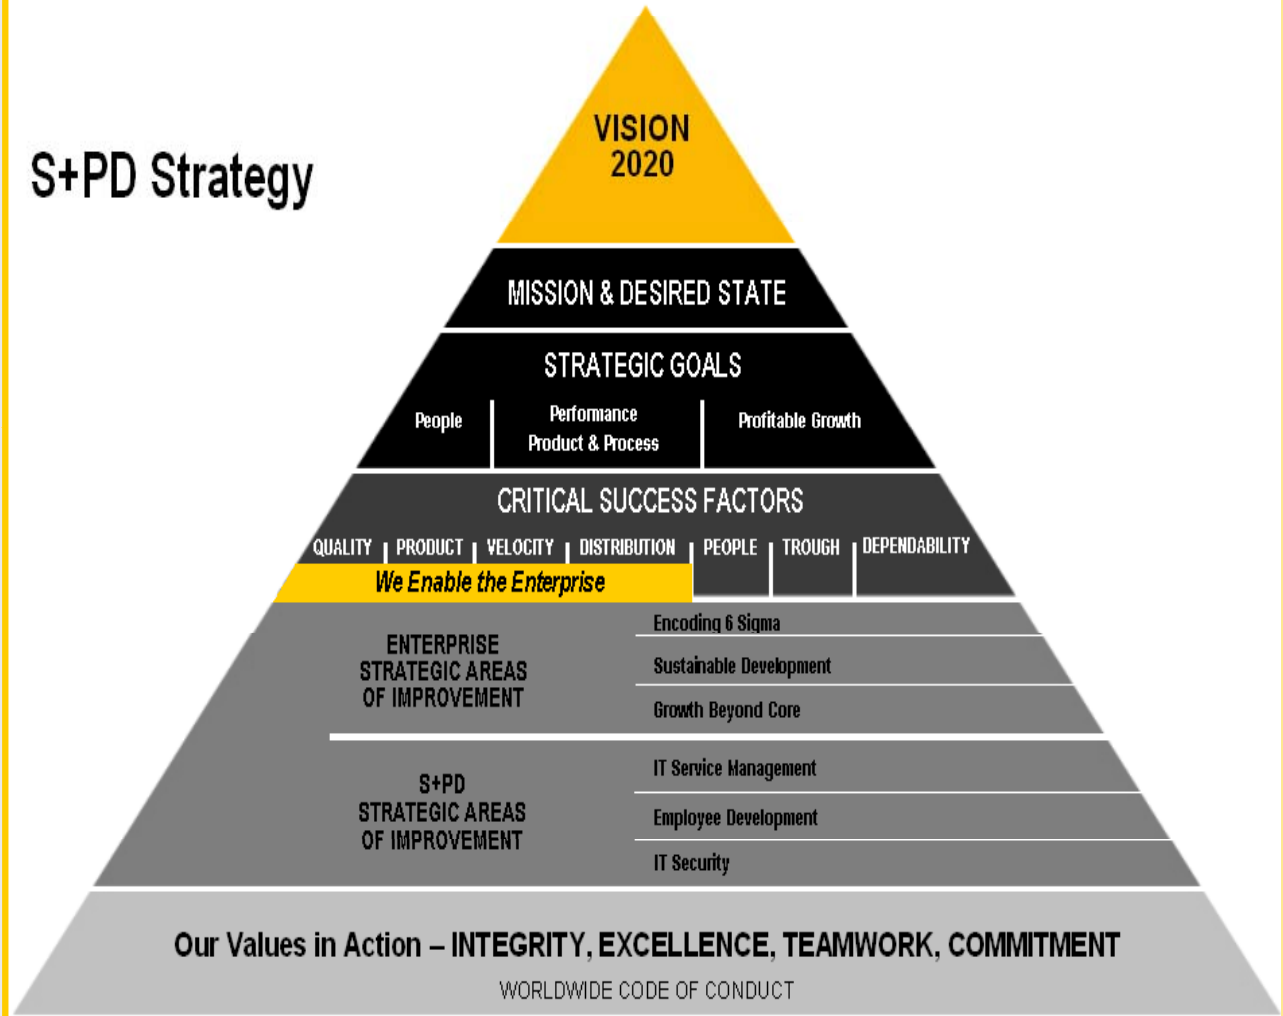
CATERPILLAR®
TODAY'S WORK. TOMORROW'S WORLD.™

The Power Behind the Progress: Information Technology

With **Vision 2020**,

Caterpillar has provided clarity of direction regarding the future of the business. It outlines what businesses we want to be in, the values we want to display, the strategic goals that will measure progress, and the critical success factors which we must focus on to realize the vision.

S+PD Strategy



Systems+Processes Division

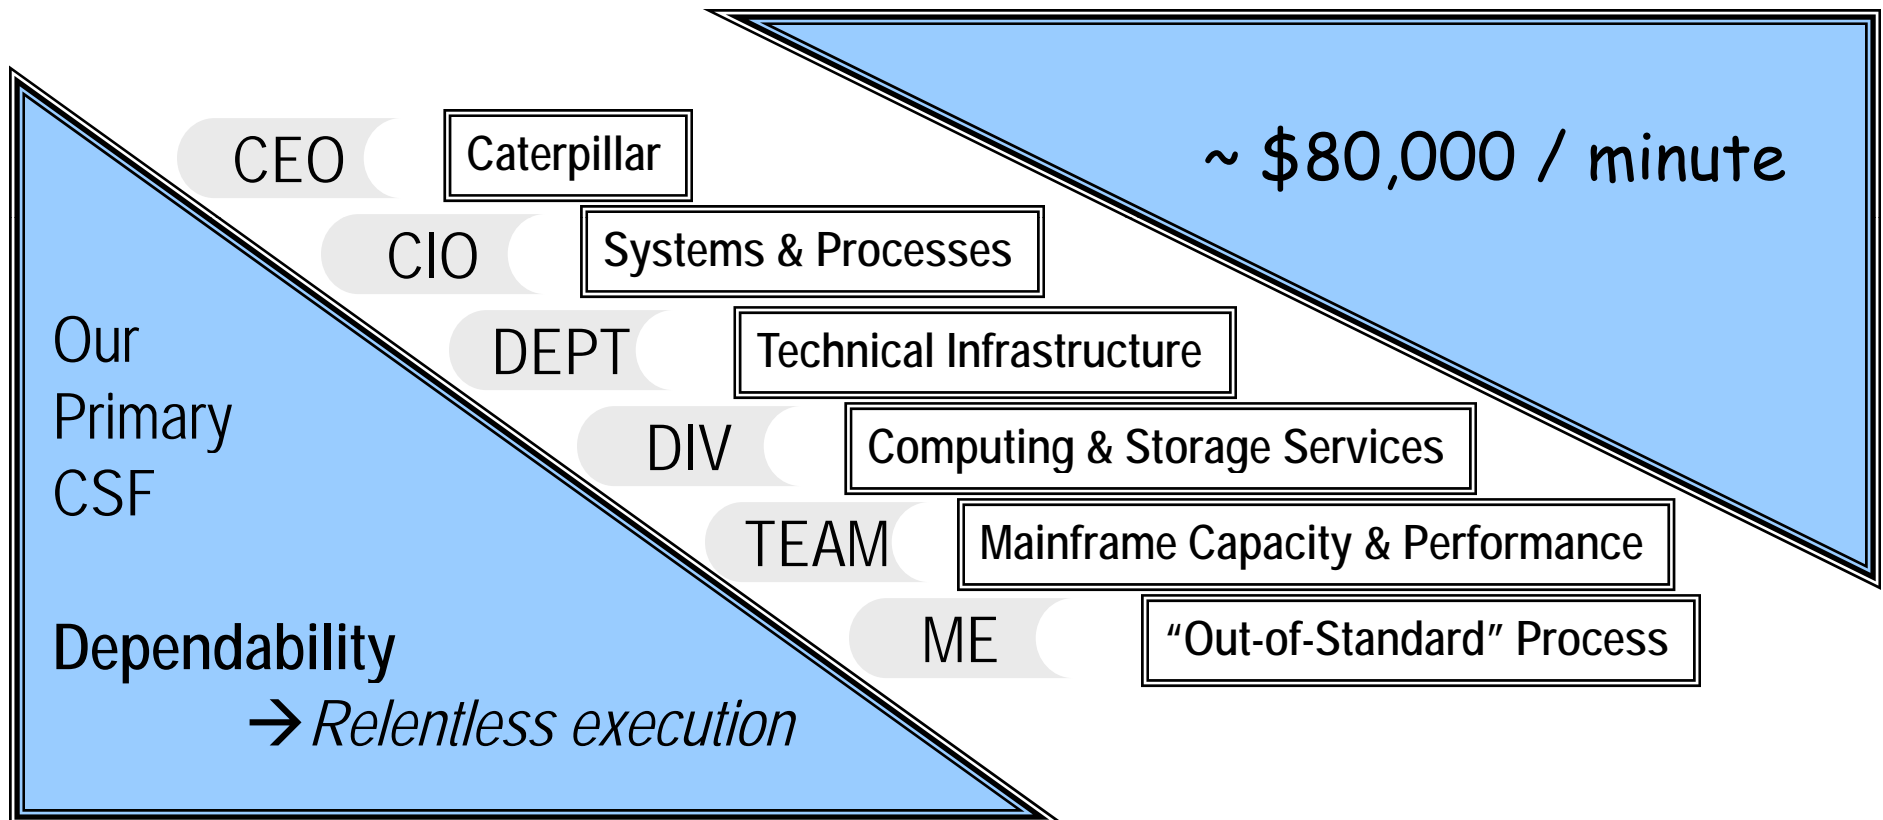
Caterpillar: Public Blue



CATERPILLAR®
TODAY'S WORK. TOMORROW'S WORLD.™

Technical Infrastructure – Computing & Storage Services

Everyday, more than 10,000 IT professionals around the world are working to enable Caterpillar's business and success.



Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

The OOS Story

Goal

Data

Process

Results

Lessons Learned & Shared

Systems+Processes Division

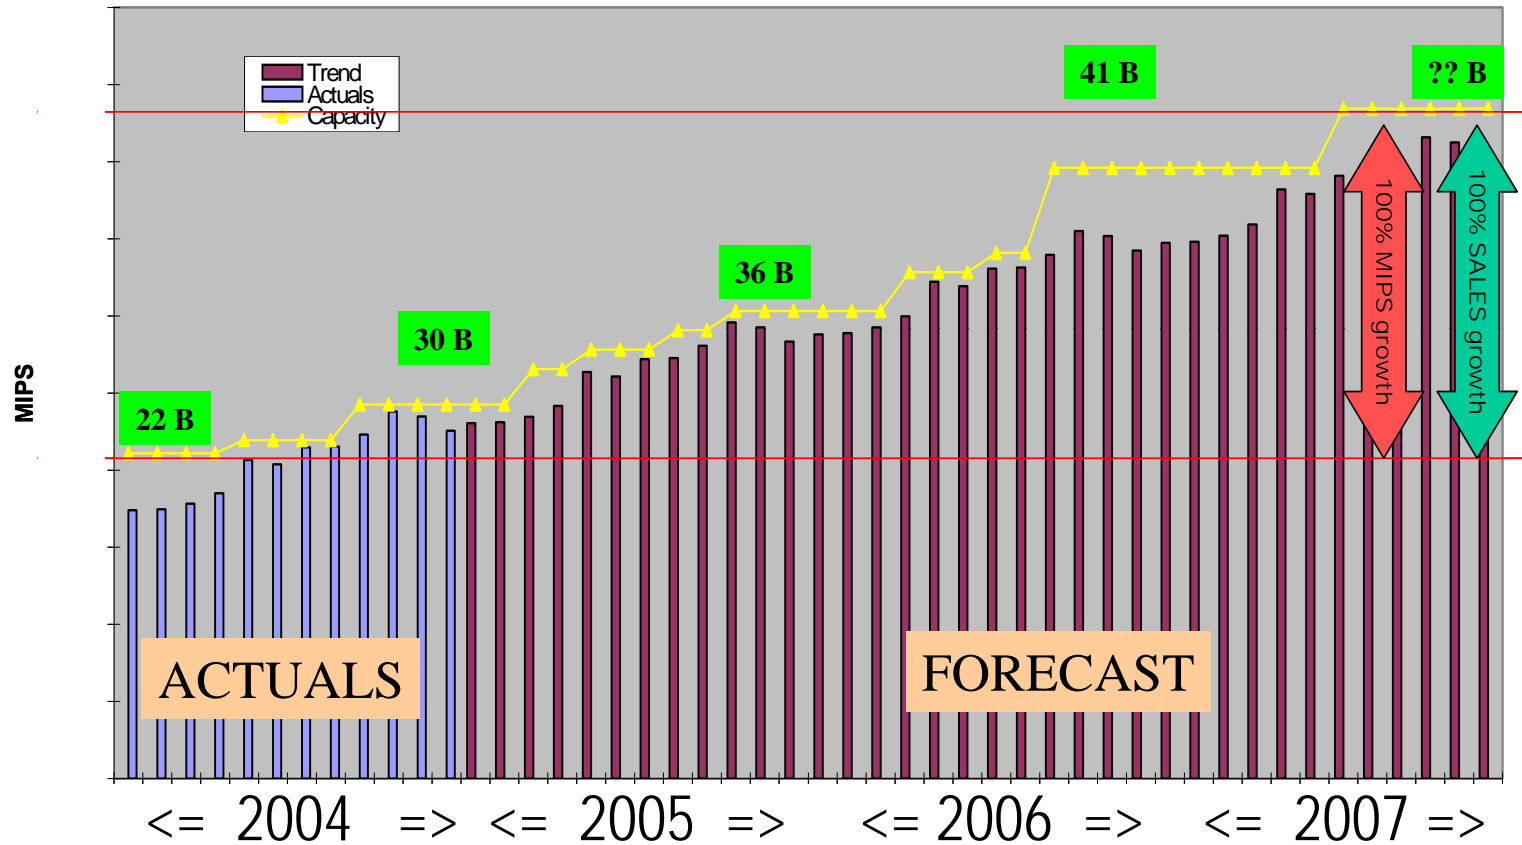
Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

GOAL of the OOS process

"to keep MIPS growth less than sales \$ growth"



Enterprise gross sales are plotted against the capacity line for mainframe growth.

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

What did we try first

- Attacking the largest applications as measured by total monthly utilization
- Questioning analyst activity based on skewed situational workloads
- Buying more capacity based on “high water mark” of utilization
- Making assessments based on average utilization and ignoring variance
- Asking the enterprise not to use the computer when we were busy

Most of these techniques were ineffective or delivered low ROI.

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Hypothesis for the OOS process

Provide an effective engine capacity management process that would:

- * establish an expected baseline for monthly application CPU utilization
- * detect “out of standard” applications by comparing monthly usage against the baseline
- * scrutinize those applications with significant variance for cause

Considering the original hypothesis, we needed three metrics in order to detect variance and investigate the cause. The first is a baseline utilization number, the second is the current application workload, and the last is the detection threshold.

At our enterprise, an application is “out of standard” when $\{(benchmark - baseline)/baseline > detection\ threshold\}$.

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

DATA requirements for the OOS process

BENCHMARK : This was the accumulation of a single month of application workload. We measured it in cpu seconds. The data begins as SMF data. After several levels of summarization, we end up with a file we think of as the daily/hourly data.

BASELINE : After collecting the benchmark data for more than 12 months, we are able to establish baseline performance for the application by simply averaging monthly totals.

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

What our BATCH data looks like

```

BROWSE      TST.Z1DW.BATCH.DAILY.G0409V00          Line 00000000 Col 001 080
Command ==>                               Scroll ==> PAGE
***** Top of Data *****
A3S3      2006-01-13  PERIOD2   AIMSARCH  AI   MS   AIMS      0.014681   2727      20
A3S3      2006-01-13  PERIOD3   AIMSARCH  AI   MS   AIMS      0.008564   2112      13
A3S3      2006-01-13  PRIME     AIMSARCH  AI   MS   AIMS      0.030800   4273      39
A3S3      2006-01-13  PERIOD3   AJVC8830  AJ   VC   AJVC      0.000083     5         1
A3S3      2006-01-13  PERIOD3   AJYY1910  AJ   YY   AJYY      0.000331    16         1
A2S2      2006-01-13  PERIOD3   AJVC8805  AJ   VC   AJVC      0.000019    63         1
A2S2      2006-01-13  PERIOD3   AJVC8810  AJ   VC   AJVC      0.000081    59         1
A2S2      2006-01-13  PERIOD3   AJVC8820  AJ   VC   AJVC      0.000083    57         1
A2S2      2006-01-13  PERIOD3   AJVC8840  AJ   VC   AJVC      0.000069    17         1
A2S2      2006-01-13  PERIOD3   AJVC8850  AJ   VC   AJVC      0.000072    16         1

```

| LPAR | DATE | SHIFT | JOBNAME | FAC | SYS | FACSYS | CPU HRS | ELAPSED | JOB COUNT |
|------|------------|---------|----------|-----|-----|--------|----------|---------|-----------|
| A3S3 | 2006-01-13 | PERIOD2 | AIMSARCH | AI | MS | AIMS | 0.014681 | 2727 | 20 |
| A3S3 | 2006-01-13 | PERIOD3 | AIMSARCH | AI | MS | AIMS | 0.008564 | 2112 | 13 |
| A3S3 | 2006-01-13 | PRIME | AIMSARCH | AI | MS | AIMS | 0.030800 | 4273 | 39 |
| A3S3 | 2006-01-13 | PERIOD3 | AJVC8830 | AJ | VC | AJVC | 0.000083 | 5 | 1 |
| A3S3 | 2006-01-13 | PERIOD3 | AJYY1910 | AJ | YY | AJYY | 0.000331 | 16 | 1 |
| A2S2 | 2006-01-13 | PERIOD3 | AJVC8805 | AJ | VC | AJVC | 0.000019 | 63 | 1 |
| A2S2 | 2006-01-13 | PERIOD3 | AJVC8810 | AJ | VC | AJVC | 0.000081 | 59 | 1 |
| A2S2 | 2006-01-13 | PERIOD3 | AJVC8820 | AJ | VC | AJVC | 0.000083 | 57 | 1 |
| A2S2 | 2006-01-13 | PERIOD3 | AJVC8840 | AJ | VC | AJVC | 0.000069 | 17 | 1 |
| A2S2 | 2006-01-13 | PERIOD3 | AJVC8850 | AJ | VC | AJVC | 0.000072 | 16 | 1 |

Each row represents a slice of workload specific to an application jobname, and delineated by LPAR, date and shift.

What our DB2 (DDF) data looks like

BROWSE MST.Z1DW.PR.DDB2HR.G0727V00 Line 00000000 Col 001 080

Command ==> Scroll ==> PAGE

***** Top of Data *****

| | | | | | | |
|------------|----------|------|----------|--------------|-----|-------|
| 2006-01-15 | 00.00.00 | Z1P2 | D1AGDMS1 | db2jccmain | 51 | 1144 |
| 2006-01-15 | 00.00.00 | Z1P6 | D1AGDMS1 | db2jccmain | 15 | 36 |
| 2006-01-15 | 00.00.00 | Z1P6 | N4GK1WEB | java | 3 | 256 |
| 2006-01-15 | 01.00.00 | Z1P2 | D1AGDMS1 | db2jccmain | 358 | 5454 |
| 2006-01-15 | 01.00.00 | Z1P9 | D1RPQRWB | G4BB508C.... | 153 | 4 |
| 2006-01-15 | 01.00.00 | Z1P3 | Z1DITRK1 | java | 69 | 37502 |
| 2006-01-15 | 01.00.00 | Z1P3 | D1AGDMS1 | db2jccmain | 55 | 4956 |
| 2006-01-15 | 01.00.00 | Z1P6 | N4GK1WEB | java | 49 | 41281 |
| 2006-01-15 | 01.00.00 | Z1P5 | D1AGTCDT | G4DE50B4.... | 36 | 2 |
| 2006-01-15 | 01.00.00 | Z1P5 | D1AGTCDT | G4DD308E.... | 33 | 2 |
| 2006-01-15 | 01.00.00 | Z1P5 | D1AGTCDT | G48B20EA.... | 32 | 2 |

| DATE | TIME | SUBSYS | AUTHID | CORRID | CPU SEC | SQL COUNT |
|------|------|--------|--------|--------|---------|-----------|
|------|------|--------|--------|--------|---------|-----------|

Each row represents a slice of workload specific to an application authid/corrid, and delineated by subsystem, date and hour.

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR
TODAY'S WORK. TOMORROW'S WORLD.™

What we found out about data

- WOW! There is a whole bunch of data out there
- Not all mainframe technologies/workloads need to be monitored
- High-level summarization can create “The illusion of average”
- It is important to know the variance of the workload as well as the average
- Summary data must be supported by the detail in order to investigate
- Monthly totals can be influenced by day-count or day-of-month pattern

Systems+Processes Division



Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Z9 demand patterns for previous year

(measured during first 5 days in PRIME only 7AM - 4PM)




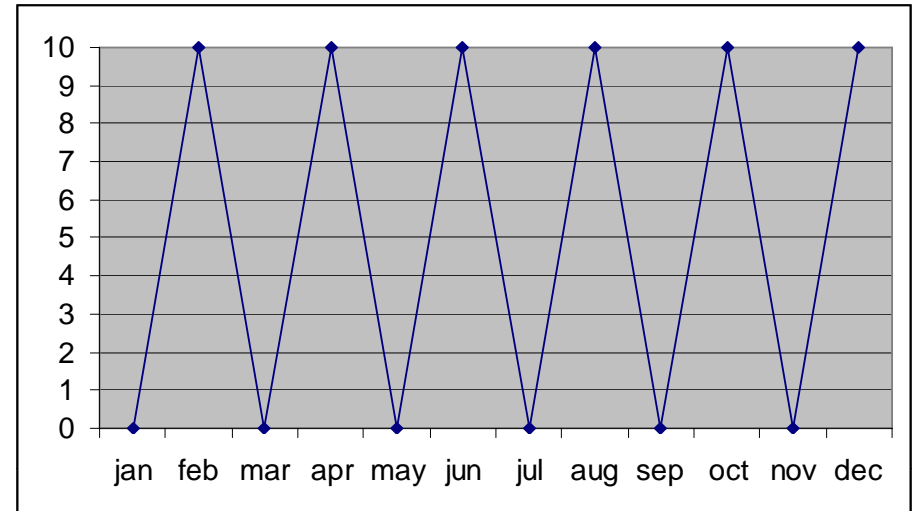

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

The ILLUSION of AVERAGE



All of these patterns average to 5. If that is all you considered, then in many cases, you may be mistaken about the true nature of a workload. This is why variance and the actual detail needs to be scrutinized.

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Day Patterns Matter

| sun | mon | tue | wed | thu | fri | sat |
|-----|-----|-----|-----|-----|-----|-----|
| | 1 | 2 | 3 | 4 | 5 | 6 |

We found that if the days in the month began like the week shown above, that the workload was harder to handle than if the month began on a weekend (below). Why? Because many workloads occur on the calendar day one regardless of whether it is a workday. When calendar day one occurs on a work day, then we have essentially a stacked workload that all must occur on the same day. This makes holidays at the beginning of a month beneficial for workload balance.

| sun | mon | tue | wed | thu | fri | sat |
|-----|-----|-----|-----|-----|-----|-----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |

Systems+Processes Division

Caterpillar: Public Blue



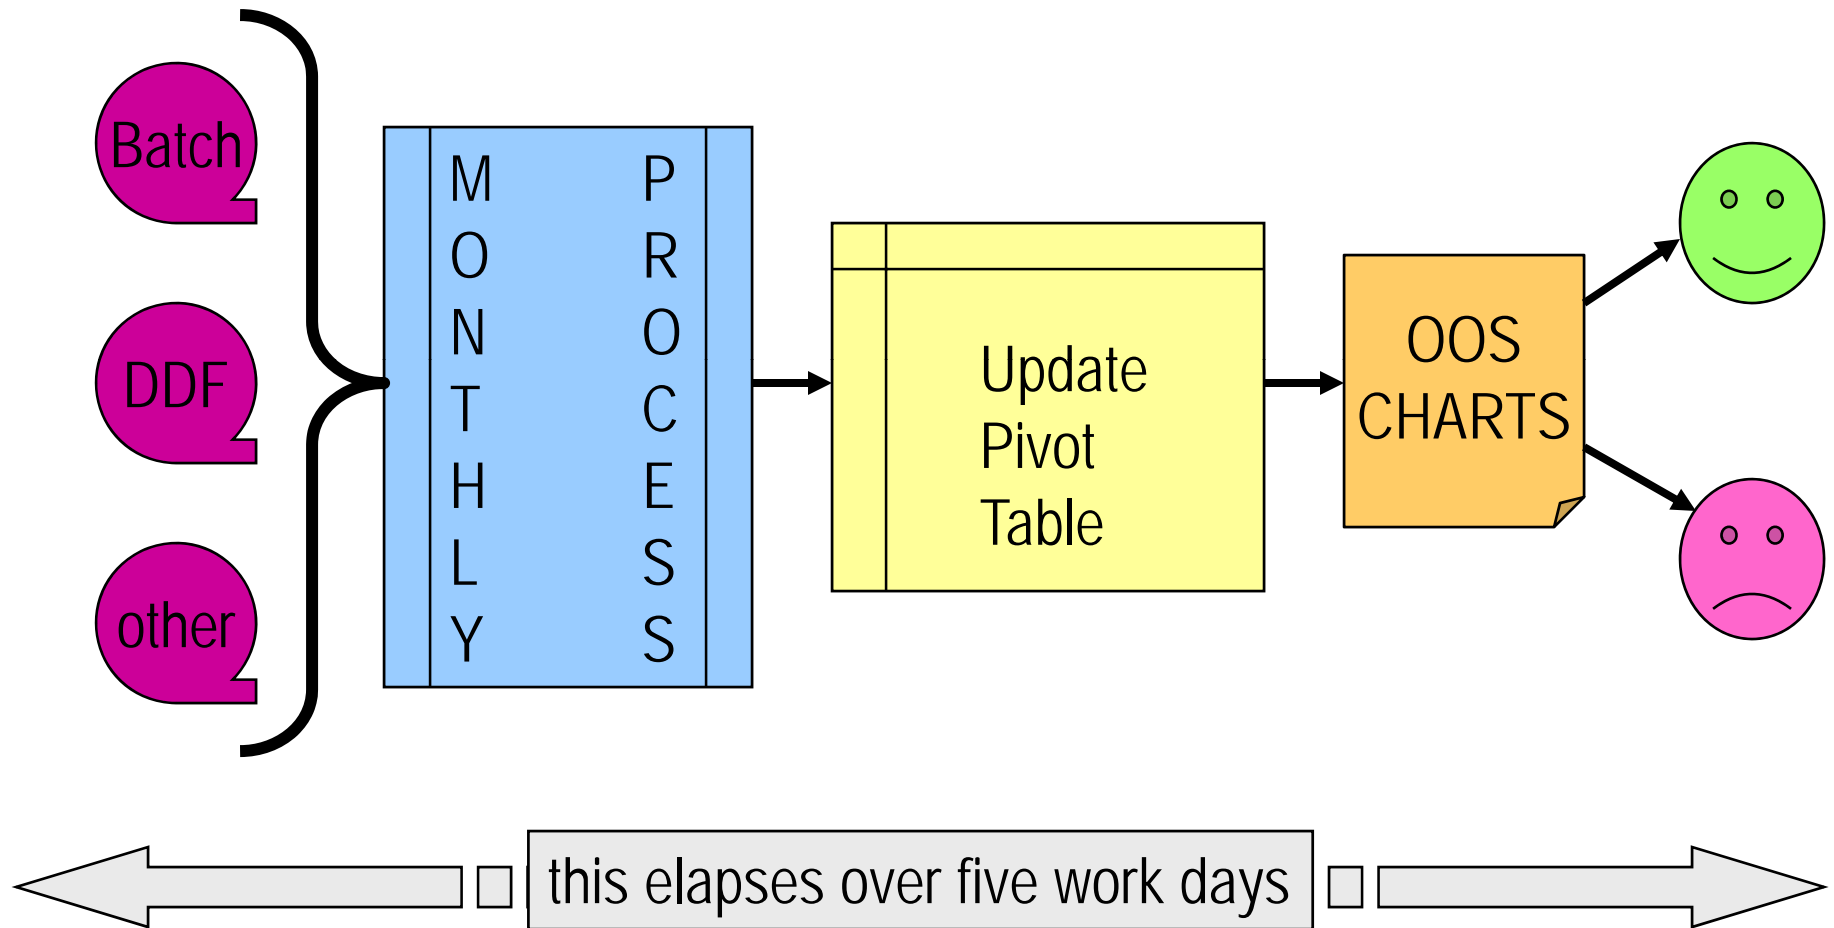
CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

The OOS process

- Start with a monthly utilization number for an application
- Collect previous history of utilization for the application
- Calculate the baseline of monthly utilization for the application
- Compare the current to the baseline
- Mark the application for scrutiny if it exceeds the detection threshold
- Determine the named cause of the variance




Automation in the OOS process



What is a “detection threshold”

Effect of Varying Detection Levels





Using a pivot table in a spreadsheet, we were able to determine how many hours of workload was out-of-standard, and the number of applications that caused the OOS. We wanted the ratio that gave us the most OOS hours from the least number of applications.



How we find the "named" cause of variance

| | | | | | | | | | |
|-------------------|-----------|-------------|-----------|---------------------|-----------|-----------|-----------|-----------|-------------|
| LPAR | (All) | % tolerance | 20% | Total systems "OOS" | 10 | | | | |
| | | min cpu se | 10,000 | Total hours "OOS" | 104 | | | | |
| Sum of normalized | DATE | | | | | | | | |
| SYS | Jul-06 | Aug-06 | Sep-06 | Oct-06 | Nov-06 | Dec-06 | Jan-07 | Aug-07 | Grand Total |
| AA | 2081 | 2440 | 2133 | 2346 | 2115 | 2079 | 2300 | 2310 | 30920 |
| AD | 18733 | 24143 | 22230 | 23833 | 23688 | 23886 | 32378 | 30116 | 362074 |
| AE | 9576 | 11490 | 10492 | 10765 | 9813 | 8016 | 9196 | 9779 | 134617 |
| AF | 442 | 494 | 447 | 506 | 1309 | 439 | 507 | 349 | 6621 |
| AG | 96 | 108 | 111 | 126 | 128 | 130 | 139 | 72 | 1346 |
| HR | 66734 | 75005 | 106945 | 57420 | 56875 | 57617 | 62708 | 89133 | 938146 |
| ZT | 31762 | 36214 | 33448 | 35171 | 32633 | 30485 | 30433 | 30505 | 435948 |
| ZZ | 1356 | 1474 | 1445 | 1437 | 1345 | 1254 | 1523 | 1487 | 19739 |
| Grand Total | 4,710,324 | 5,120,942 | 4,956,164 | 4,821,923 | 4,910,725 | 4,492,532 | 5,179,022 | 5,499,338 | 71,020,847 |



Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR
TODAY'S WORK. TOMORROW'S WORLD.™

How we find the "named" cause of variance

| | | | | | | | | | | | | | | | | | | |
|-----|------|------------|---------|----------|------|------|------|---------|------|-----------------|------|------|------|------|------|------|------|------|
| 001 | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ANOMALY PATTERN | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- |
| 002 | CASA | 2007-08-01 | PRIME | Z1HRPSA7 | Z1 | HR | Z1HR | 13.0222 | | | | | | | | | | |
| 003 | CASA | 2007-04-01 | PRIME | Z1HRPSA7 | Z1 | HR | Z1HR | 0.7367 | | | | | | | | | | |
| 004 | CASA | 2007-02-01 | PRIME | Z1HRPSA7 | Z1 | HR | Z1HR | 1.0690 | | | | | | | | | | |
| 005 | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- | GROWTH PATTERN | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- |
| 006 | A1S1 | 2007-08-01 | WEEKEND | Z1HR1900 | Z1 | HR | Z1HR | 0.8123 | | | | | | | | | | |
| 007 | A1S1 | 2007-04-01 | WEEKEND | Z1HR1900 | Z1 | HR | Z1HR | 0.7456 | | | | | | | | | | |
| 008 | A1S1 | 2007-02-01 | WEEKEND | Z1HR1900 | Z1 | HR | Z1HR | 0.6789 | | | | | | | | | | |
| 009 | A1S1 | 2007-08-01 | PRIME | Z1HR1810 | Z1 | HR | Z1HR | 0.5751 | | | | | | | | | | |
| 010 | A1S1 | 2007-04-01 | PRIME | Z1HR1810 | Z1 | HR | Z1HR | 0.3351 | | | | | | | | | | |
| 011 | A1S1 | 2007-02-01 | PRIME | Z1HR1810 | Z1 | HR | Z1HR | 0.2751 | | | | | | | | | | |
| 012 | A1S1 | 2007-08-01 | PERIOD2 | Z1HR5010 | Z1 | HR | Z1HR | 3.1266 | | | | | | | | | | |
| 013 | A1S1 | 2007-04-01 | PERIOD2 | Z1HR5010 | Z1 | HR | Z1HR | 2.4787 | | | | | | | | | | |
| 014 | A1S1 | 2007-02-01 | PERIOD2 | Z1HR5010 | Z1 | HR | Z1HR | 1.2387 | | | | | | | | | | |
| 015 | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- | TUNING PATTERN | ---- | ---- | ---- | ---- | ---- | ---- | ---- | ---- |
| 016 | A1S1 | 2007-08-01 | PERIOD2 | Z1HR2064 | Z1 | HR | Z1HR | 5.9884 | | | | | | | | | | |
| 017 | A1S1 | 2007-04-01 | PERIOD2 | Z1HR2064 | Z1 | HR | Z1HR | 2.4234 | | | | | | | | | | |
| 018 | A1S1 | 2007-02-01 | PERIOD2 | Z1HR2064 | Z1 | HR | Z1HR | 2.4566 | | | | | | | | | | |

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Patterns of utilization variance

- Anomaly → This is a workload that will not occur on a permanent monthly basis in the future. We investigate cause to prevent the workload if possible, or determine the cyclic nature of the workload for future forecasting.
- Growth → This is a workload that will add into your enterprise baseline every month. It is useful to understand the size and frequency of this workload. We investigate to try and ensure standards of performance are met.
- Tuning → This is a workload that we believe to be inefficient based on expected standards of performance. Tuning can be as simple as frequency adjustment, or data maintenance. More complex tuning involves adjusting the “unit-of-work”.



RESULTS from the OOS process

- We discovered a queue of tuning opportunities
- We were able to identify workload anomalies to control/ignore
- We identified growth in workloads and load balanced
- After 3-6 months, application variance began to reduce
- Performance stabilized in almost all enterprise applications
- We began to develop a stable workload pattern in PRIME
- Reactive and Proactive capacity increases have been controlled

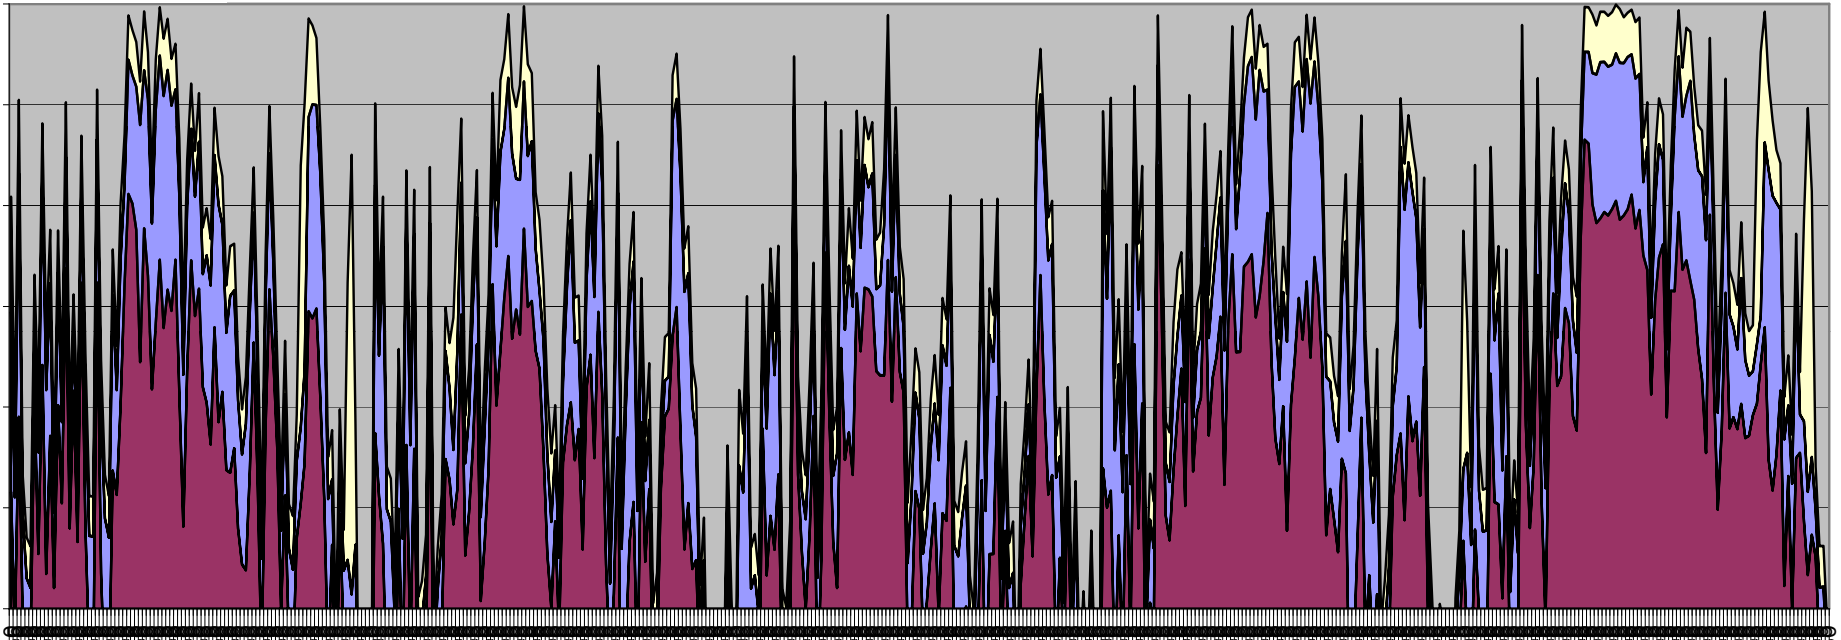
Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

A better utilization/capacity picture



Reduced workloads

Improved Load
Balancing

Less variability

Systems+Processes Division


Caterpillar: Public Blue

RESULT



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Control of the MIPS growth rate - accomplished



Systems+Processes Division

Caterpillar: Public Blue



LESSONS LEARNED from the OOS process

Capacity & Performance Wisdom
we gathered during the OOS process

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Workload source & What you can affect

Σ (freq * unit-of-work)

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Don't waste time on
workloads that you
cannot change.

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

WHEN a workload runs
can be more important
than the workload size.

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

3 methods to control capacity

EMT

Eliminate workloads

Move workloads

Tune workloads



Three types of support staff that you will meet.

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Delineation of
functionality
creates better
Diagnostics.

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Sizing workloads
can improve
application
forecasts.

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Pad management
protects
performance.

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Standardize your
notifications & expectations.

NOT the responses.

Systems+Processes Division

Caterpillar: Public Blue

LESSONS



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™

Thank you

Any Questions ???

CAT, CATERPILLAR, their respective logos, “Caterpillar Yellow” and the POWER EDGE trade dress, as well as corporate and product identity used herein, are trademarks of Caterpillar and may not be used without permission.

© 2008 John Van Wagenen, Caterpillar Inc.

Systems+Processes Division

Caterpillar: Public Blue



CATERPILLAR[®]
TODAY'S WORK. TOMORROW'S WORLD.™