Honeywell

Determine A System's Performance Characteristics

L61531
LCN

TotalPlant

Notices and Trademarks

This module supports **TotalPlant** Solution (TPS) system network.

TPS is the evolution of TDC 3000X.

Honeywell Inc.

Industrial Automation and Control

Automation College

2820 West Kelton Lane

Phoenix, AZ 85053-3028

1-800 852-3211

Table of Contents

Figures

Tables

Acronyms, Parameters, and Resources Referenced Later In This Module

Acronyms

ACIDP..Advanced Control Interface Data Point
ACP ....Advanced Control Program
CRDP ....Calculated Results Data Point
HCU...Honeywell Control Unit
LSTF...Large System Test Facility
PIN...Plant Information network
PSDP...Processor Status Data Point
SMCC ....System Maintenance Control Center
ULP...Upper Level Processor

Resources

LCN Guidelines – Implementation, Troubleshooting, and Service	*LC09-510*	*(TPS 3025)*
Engineer's Reference Manual	*SW09-605*	*(TPS 3030-2)*

MODULE INTRODUCTION

To evaluate system performance, we need a baseline or some sort of relative measurement. For this purpose, Honeywell developed the “cluster and “Honeywell control unit (HCU)” concept as a basis from which to measure TPS Network system performance and integrity. The cluster model and the HCU make up an artificial configuration to provide a consistent system whose characteristics Honeywell can measure from release to release.

The Processor Status Data Point (PSDP) provides access to the operational characteristics of all LCN nodes. In this course module, you will look at the PSDP information available on a system, determine the system loading, and, through a process called “baselining,” determine the operational characteristics of the system during normal operation.

Before performing software or hardware upgrades or making an operational change, it is important to know the normal operational characteristics of the system in order to make comparisons later to determine the effects of the change on the system.

Objective

Determine the performance of a TPS Network LCN system by interpreting the performance parameters of the LCN nodes.

Performance Testing

The Cluster Model

Because the LCN is a complex network of different types of nodes and the TPS Network system is a flexible system that can exist in many different configurations to satisfy specific customer applications, it is not feasible to test all possible configurations and network loads. In an effort to understand the effect of network configuration and the impact of new software releases on system performance and integrity, the performance cluster was developed. This is a performance testing mechanism used to validate generic performance of the system and to estimate performance in untested configurations.

An operator in a plant is usually responsible for a specific process that employs a set of equipment to produce a product. The performance cluster concept aligns the LCN-based TPS Network equipment and Area needed to support such an operator. See Figures 1 and 2.

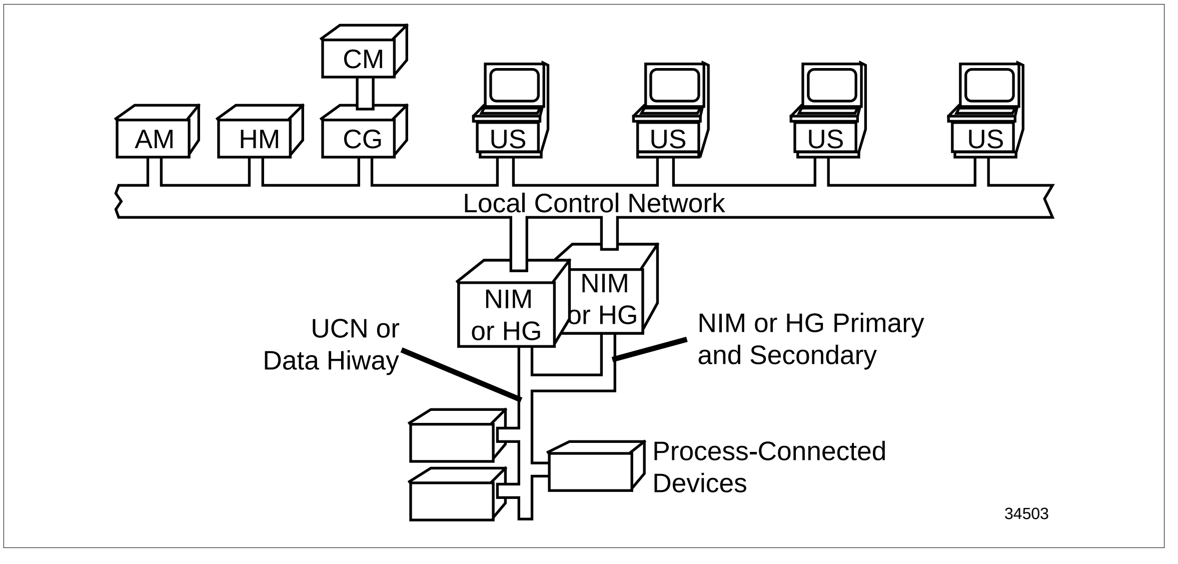

Figure 1 - The Performance Cluster

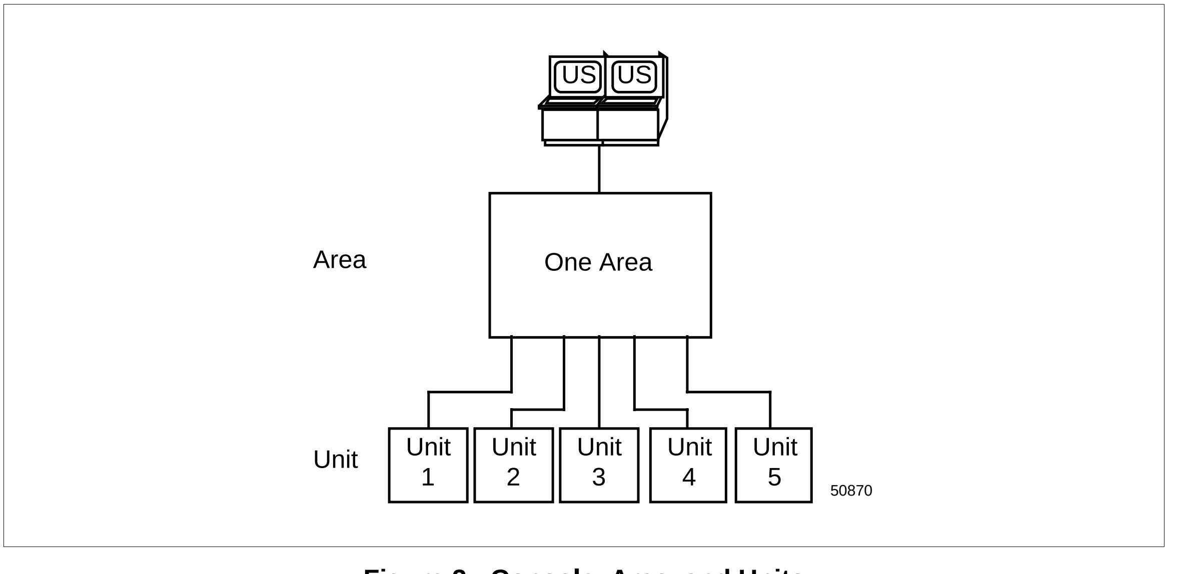

Figure 2 - Console, Area, and Units

What Is A Cluster?

Functions

A single performance cluster represents these functions:

- storing historical data and system software,
- performing advanced process control,
- connecting the LCN to the process-connected hiways/networks,
- providing human interfaces.

Hardware

A single performance cluster of nodes supports one process gateway and consists of the following hardware:

- 1 - console with 4 USs (one Area with one or more physical processing Units),
- 1 - Data Hiway with process-connected devices,
  or
  1 - UCN with process-connected devices,
- 1 - HM (journals, history, checkpointing, Area database),
- 1 - AM (advanced control),
- 1 - CM (used by multiple clusters).

In actual practice, the CM is used across several clusters, but for performance testing it is assigned to one cluster, because testing a “fractional” CM is not practical.

Multiple Clusters

For multiple clusters, Honeywell recommends multiple HMs:

- One HM to support cluster-specific activities, such as history, journals, checkpoints, Area databases, and schematics.
- One additional HM to store system-related data and optional user volumes across all clusters.

How Are Clusters Used?

The definition of a cluster encompasses a large percentage of the installed TPS Network systems. Any system can be approximated by using multiple clusters, and the performance of the system can be estimated based on the measured performance of one cluster.

ATTENTION

ATTENTION—The performance values given in this course module apply to only the cluster with the mix of equipment described here. On a very large system, it is recommended that you partition as closely as possible to the cluster model. The basic idea is to distribute the load in order to maximize performance. If you do not follow the cluster model, this does not mean that you will necessarily have poor performance; however, performance will be improved if the cluster model is followed.

The Honeywell Control Unit

After defining the basic LCN node arrangement with the cluster model, we need to define the activity level of the system whose performance we are measuring. The Honeywell Control Unit (HCU) is the second concept upon which performance measurements are based. This is a database that is a subset of a typical process. It is a mix of point types representative of a generic 10-valve control process. The HCU can be multiplied to simulate larger control system loads for performance studies.

Each cluster supports 10 to 20 HCUs.

REFERENCE—Refer to the *Engineer's Reference Manual* for the mix of points in an HCU.

Factors that Affect System Performance

Configuration Guidelines

The following are some guidelines that should be met when configuring your system:

General:

- Follow the cluster concept as closely as possible. A cluster should handle control for just one area.
- Take care that cross-cluster activity does not become a normal operating condition because this will increase the load on the NIM/HG and the AM.

Configuration:

- Reduce the complexity of schematics. Complex schematics are less efficient to display and take longer to call up to the screen.

Hardware:

- Both NIM/HG in a NIM/HG pair should be on the same LCN segment. They should not be separated by a fiber optic link.
- All US/GUSs in the same console should be on the same LCN segment. The NIM/HG with which the Universal Station (US)/GUS Universal Station (GUS)s often communicate should also be on the same segment.
- The two clock-source nodes should be on the main segment.
- Each LCNE can be housed in the chassis for one of the following node types: NIM/HG, AM, HM, CG or US. Because the NIM/HG is the least likely to be turned off, they should be considered first.

Example Symptoms of a "Loaded" System

You know you have a system-loading problem when one or more of the following symptoms is reported:

- AM overruns are occurring, and you know the AM may soon crash.
- Slow display callup and response to operator actions. The operator is nervous; he may lose control of the process.
- The Event Recovery target keeps appearing. If a plant upset occurred, consistent alarming may be hampered.
- The message" HM nn RECNFG HM HST OVRLD HISTORY MODULE" has been journaled and now history gaps exist.
- A Journal Buffer Overflow message printed on the printer and now the History Module must be the source for operational history.

At the time these problems are seen, this system is loaded. These examples are not inclusive, but are examples of problems that occur when a system is pushed to its limit.

Using PSDPs to Measure TPS Network System Performance

Processor Status Data Point

Definition

A Processor Status Data Point (PSDP) is an internal point established on the LCN when a node is loaded. This dedicated internal point is referenced by this name:

`$PRSTSnn.parameter`

where:
`nn` = the node number (01-64)
`parameter` = the specific parameter name for the desired data.

Characteristics

Some general characteristics of the PSDP are

- They are available in every loaded node on the LCN.

- A portion of the PSDP parameters are supplied at the software environment (SE) level:
  – CPU and HEAP values,
  – certain node and cable status values,
  – data access related parameters.

- The remaining PSDP parameters are node-specific and supply data applicable to the node type:
  – processing rates,
  – memory utilization,
  – hardware interface information.

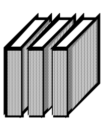

REFERENCE—The *Engineer's Reference Manual* section 22 describes all of the PSDP parameters designed for customer use.

Additional reference information is also available in the LCN *Guidelines* manual.

PSDP Parameters Concerning LCN Read/Write Access

PARSECS, RPARSECS, and SPARSECS

Figure 3 illustrates five PSDP parameters that are used to measure node performance. Tables 1 and 2 describe the parameters.

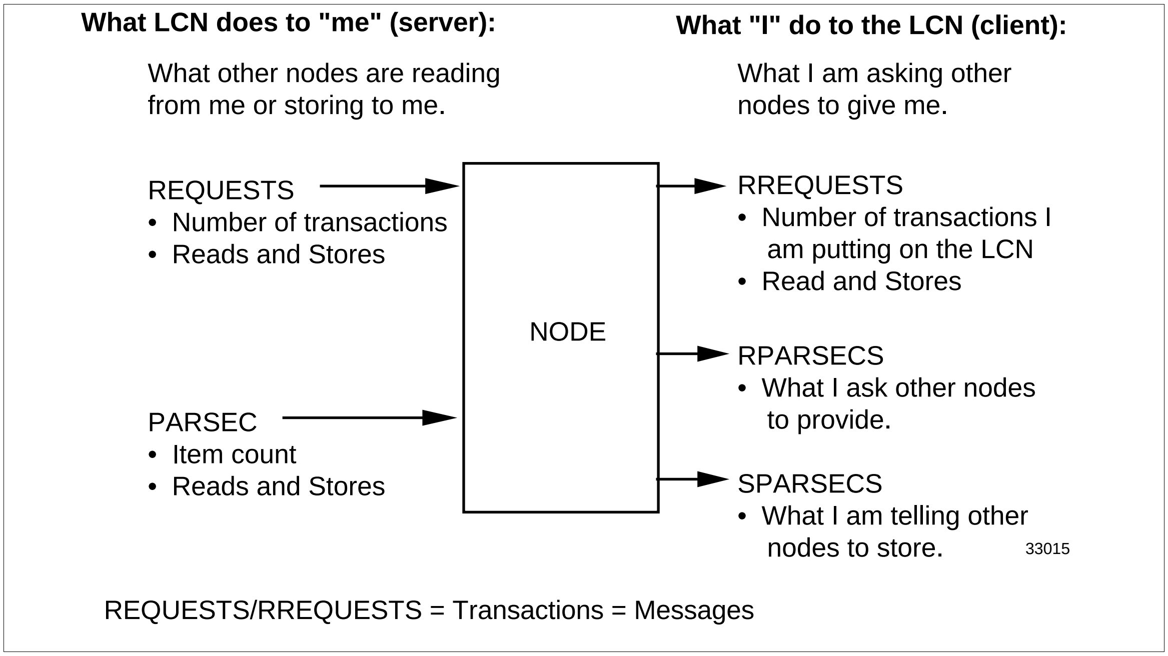

Figure 3 - Diagram of Various PARSEC Parameters

Table 1 - Description of PARSEC, RPARSEC, and SPARSEC

Parameter	Description
REQUESTS	Measures the number of transactions (messages) involved with other nodes reading from "me" or storing to "me."
PARSEC	Measures the number of parameters that other LCN nodes are asking me to provide (through reads from me), or are telling me to store (through writes to me). This is an item count. Is usually comprised of mostly reads—what this node is supplying to everybody else.
RREQUEST	Measures the number of requests I am putting out on the LCN for either reads or stores. Measures what I am doing on the LCN.
RPARSEC*	What I am asking other nodes to provide to me through reads from the LCN.
SPARSEC*	What I am telling other nodes to store through writes to the LCN.

*RPARSEC and SPARSEC can provide help in determining what nodes are causing too large a PARSEC value on some other node on the LCN.

Table 2 - Node Comparison of PARSEC, RPARSEC, and SPAR

Node Type	Parameter
US/GUS	REQUESTS and PARSEC—Small. Measures requests coming into the US/GUS from other LCN nodes. Includes console and Area Manager-type functions. RPARSEC—Large. This measures what the US/GUS is asking other nodes to provide it. SPARSEC—Small. This measures what the US/GUS is telling other nodes to store (such as through targets and operator SP changes).
AM	REQUESTS and PARSEC—Probably fairly small. Depends on whether CGs or other AMs fetch and store data from this AM. Depends on whether US/GUSs need AM data for Schematics and Group displays. RPARSEC—Varies. Measures what the AM is asking other nodes to provide it (prefetches). SPARSEC—Varies. Measures what the AM is storing to other nodes such as HGs and NIMs (post stores). This will depend on the control strategies.
HG/NIM	REQUESTS and PARSEC—Large. The data owners are being called on to provide a large amount of information to the LCN (US/GUSs, AMs, CGs, HMs), or are being asked to store data (US/GUSs, AMs, CGs). RREQUESTS, RPARSEC, SPARSEC—Zero. The true data owners do not make any requests on the LCN.

Example Request

One message with 100 reads and 20 stores translates into

- 1 Request
- 120 PARSEC (even though node transmits back 100 parameters)

Example AM Control Strategy

Assume you have the following control strategy:

- The CG provides information to the AM.
- The AM then writes to the NIM or HG.
- The US uses data from the AM in Schematics and Group displays.

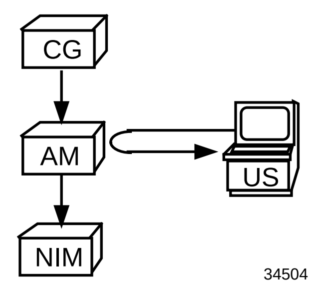

Figure 4 - Control Strategy Example

In this above control example, the AM performance parameters would indicate the following values:

Parameter	Value
REQUESTS and PARSEC	May be large, depending on the number of transactions and the number of parameters per message that other nodes are reading from or storing to the AM. There may be a large number of stores from the CG to the AM. The US will also contribute to these counts for the AM for graphic support.
RREQUEST	May be large, depending on the number of stores from the AM to the NIM or HG.
RPARSEC	Zero. The AM is not to obtaining data from the NIM or HG.
SPARSEC	This will always be some number because the AM is storing data to the NIM or HG.

The Most Helpful PSDP Parameters

Table 3 describes the PSDP parameters that are the most helpful in determining the operational characteristics of each node on the system.

Table 3 - PSDP Parameters—The Most Helpful

<table>
<tr><th>PSDP Parameter</th><th>Description</th></tr>
<tr><td>HEAPFREE[1,2]</td><td>The amount of heap available in the node at a given time.

Should be in the 100,000+ range.

Experience has shown that the AM HEAPFREE may be at very small value (20,000 - 30,000) and still operate properly.</td></tr>
<tr><td>HEAPMIN[1,2]</td><td>The minimum amount of heap the node has seen since it was started.

Should be above 75,000, except for the AM.</td></tr>
<tr><td>HEAPFRAG[1,2]</td><td>Heap fragmentation value.

Values range from 0 to 100:
• Usually 0 or 1 (0 = no fragmentation)
• Values above 3 or 4 are rare and indicate a potentially fatal condition.
• 100 = 100% fragmentation.

When heap is taken from the pile, the amount of heap needed must be consecutive for a given task. When the task is completed the heap is placed back in the same place.

Any sustained fragmentation of heap should be investigated.

It is not known what value of HEAPFRAG will actually cause a node to crash. At times, a value of 3 could cause a node to crash; however, in the case of a Universal Station, reports have listed HEAPFRAG as high as 6, without crashing the node.

NOTE: US/GUSs running engineering functions (such as DEB, CL, and PE) will typically show a HEAPFRAG of 1 to 3, but should return to 0 when the engineering function is terminated.</td></tr>
<tr><td colspan="2">NOTES:

1. Heap is a portion of static memory that is used by the CPU to perform a given task. After that task is complete, the heap memory is released for use by another task, keeping the amount of memory required by each node to a minimum.

2. All HEAP values are in 16-bit words.</td></tr>
</table>

Table 3 - PSDP Parameters—The Most Helpful, continued

PSDP Parameter	Description
CPUFREE	This is the percentage of free time the CPU has averaged over a 15-second period. As the amount of task the CPU is asked to perform increases, the amount of CPU free time decreases. Monitoring this parameter in conjunction with the parameters per second (PARSEC) gives you a good indication of the node's performance. Should not be below 15 - 20%, except for short periods (seconds). Should average •30%. NOTE: 68000—will be significantly less than the 68020 processor in the same loading scenario. 68020/68040—CPUFREE values are not accurate, because of cache memory effects. Nodes may never show a value below about 5%, so consider 5% as "near zero."
PARSEC	Indicates the number of parameters per second other nodes are reading from or storing to this node. Should not exceed the rating for that node type. Some typical ratings are • 68020 HG 630 PAR/SEC + 90 PT/SEC AM LOAD • 68040 HG 960 PAR/SEC + 120 PT/SEC AM LOAD • 68020 NIM 750 PAR/SEC + 90 PT/SEC AM LOAD • 68040 NIM 1400 PAR/SEC + 120 PT/SEC AM LOAD • 68020 AM 90 PT/SEC + 410 PAR/SEC SUPPORT • 68040 AM 120 PT/SEC + 615 PAR/SEC SUPPORT CAUTION: This count can be misleading because certain parameter requests count as 1, but translate into a heavy load on the data owner. Examples are full array accesses and Group and Detail displays. The Group display uses a "collection set" for each point, which counts as only 1 parameter in a display update, but actually translates into 15 - 20 parameters per second in the data owner.

Node-Specific PSDP Parameters

The following tables list, by node type, some of the more important PSDP parameters. There are many more than those listed here.

AM-Specific PSDP Parameters

The AM has over 100 PSDP parameters available to provide information on memory utilization, operation, loading, and optional settings; Table 4 lists those of most critical interest.

Table 4 - PSDP Parameters—AM

PSDP Parameter	Description
PRAVGx*	Indicates the point/sec load being processed by the AM.
PFAVGx*	Indicates the prefetch parameter read rate.
PSAVGx*	Indicates the poststore parameter store rate.
ALBURST	Indicates the burst alarm rate since last reset.
PFPSOVRC PFPSOVRP	Indicates the number of times the prefetch/poststore within a cycle was up to 2.5 seconds behind schedule.
PFPSHLDC PFPSHLDP	Indicates the number of times control has had to “hold its breath” for the off-node data requests to catch up, because the prefetch/poststore was more than 2.5 seconds behind schedule.
PFPSOVER	Indicates the number of cycles the off-node data prefetch/poststore request has been returned late or behind.
OVERRUNS	Indicates the number of fast processor cycles in which control has fallen behind.
AMOVRABT	AM is set to crash if control falls this far behind. Default = 100 cycles. If desired, this parameter can be set to a value other than the default (entering a -1 value prevents the AM from crashing because of falling behind).

* Those parameters whose name ends with “x” have three versions available where “x” is

S - The value over the specified snapshot period, typically 10 seconds.

C - The value as accumulated during the current hour.

P - The value as accumulated for a full 60 minutes during the previous hour, except for the hour when the node was loaded.

Table 5 - PSDP Parameters—HG and NIM

PSDP Parameter	Description
EVTRATE	The 15-second count of events, alarms, and returns-to-normal, which occurred in the last check period.
DHPARSEC(i) DHMSGSEC(i) i = index number 1-5	Represents the data access loading requests placed on the NIM or HG. 1 - regulatory control requests from an operator 2 - Node requests from AM, CG, HM (NIM only) 3 - HM requests (HG only) and display invocation 4 - Display updates 5 - Engineering functions

Table 6 - PSDP Parameters—HM

PSDP Parameter	Description
NMCHPTS	Gives the number of actual parameters in continuous history.
CHCYCAV	The average history collection cycle. For 60-second history, the average should probably be <35 seconds to ensure the maximum is never exceeded, which can lead to the infamous “HST OVERLOAD” message.
CHCYCMX	The maximum history collection cycle time in seconds per type of history seen in the last hour. This value should never exceed the time in the collection period (5, 10, 20, 60 seconds) or the HM may be getting close to having overruns and history gaps.

Table 7 - PSDP Parameters—CG

PSDP Parameter	Description
XMITRAT	The number of words transmitted to the Upper Level Processor (ULP) in the last minute.
RECVRAT	The number of words received by the CG from the ULP in the last minute.
GPARGRT	The number of parameters read from LCN nodes in the last minute.
GPARSRT	The number of parameters stored to LCN nodes in the last minute.
GHISTRT	This is the “Get History” request rate from the ULP.

Table 8 - PSDP Parameters—US/GUS

PSDP Parameter	Description
NMSCHPR	The number of parameters requested in the last custom display called by the Universal Station.
AREAALM	The last 15-second count of events, alarms and returns-to-normal seen by the US.
CURRDSP	Specifies which display is on the US/GUS screen. If the display is a picture editor schematic, it shows the filename of the schematic. If the display is a standard display, it shows the name of the display

Performance Displays

Standard Displays

Standard displays are provided with your system (on the &C6 cartridge in the DIA1 and TLK1 directories) that contain menu-driven schematics to access some of the PSDP parameters. You should try to determine your system load only when you are operating under normal conditions, allowing you to determine the following:

- Does the system has room for expansion?
- Does the system have room for unexpected upsets?

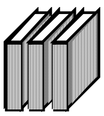

REFERENCE—The PSDP parameters are documented in the *Engineer's Reference Manual, Section 22.* Use the manual to help you determine if your system is reaching any system limits.

Performance Display Menu

Figures 5 and 6 show the pages of the PERFMENU display (menu of performance and loading displays). The menu contains targets to access the various Honeywell-supplied displays. Beside each target is a brief description of the display. There is a version number shown in the upper right corner of each display representing the latest software release that a change was made in that display. In some cases the version shown in the display will be different than the version of software currently being used. See Table 9 for what version each display should show for each running release of software.

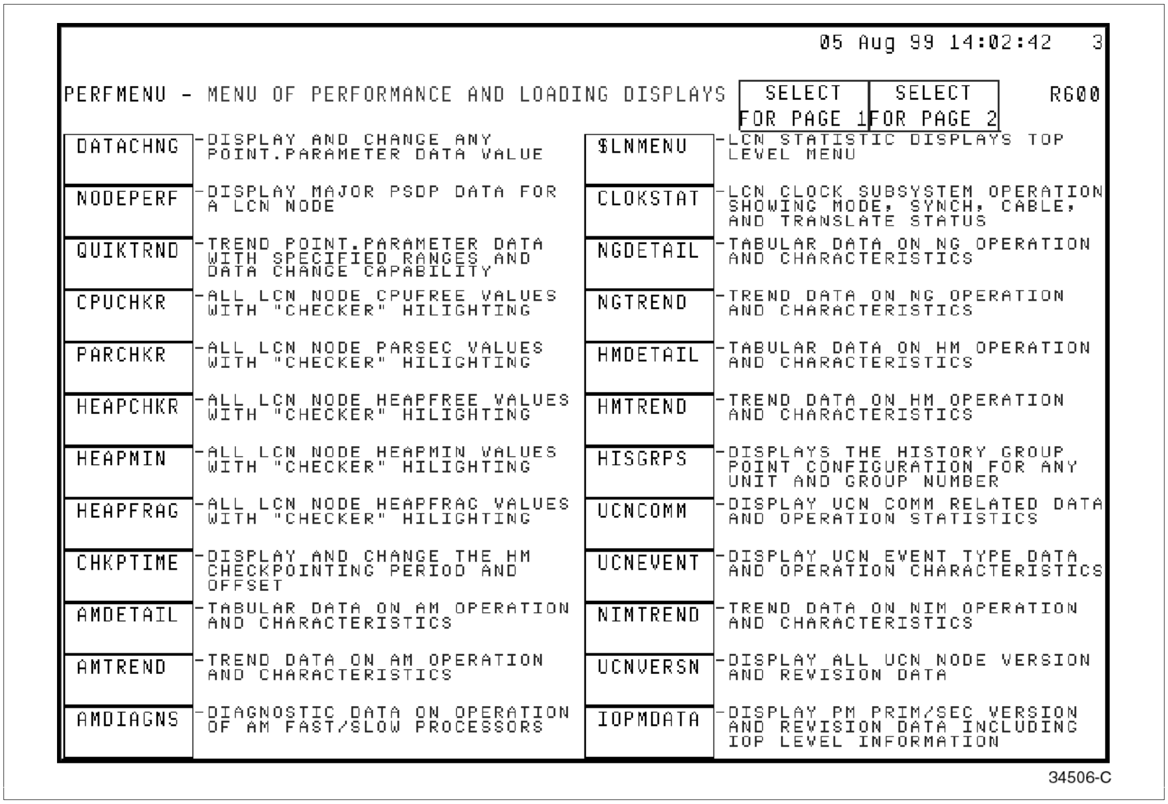

Figure 5 - Menu of Performance and Loading Displays—Page 1

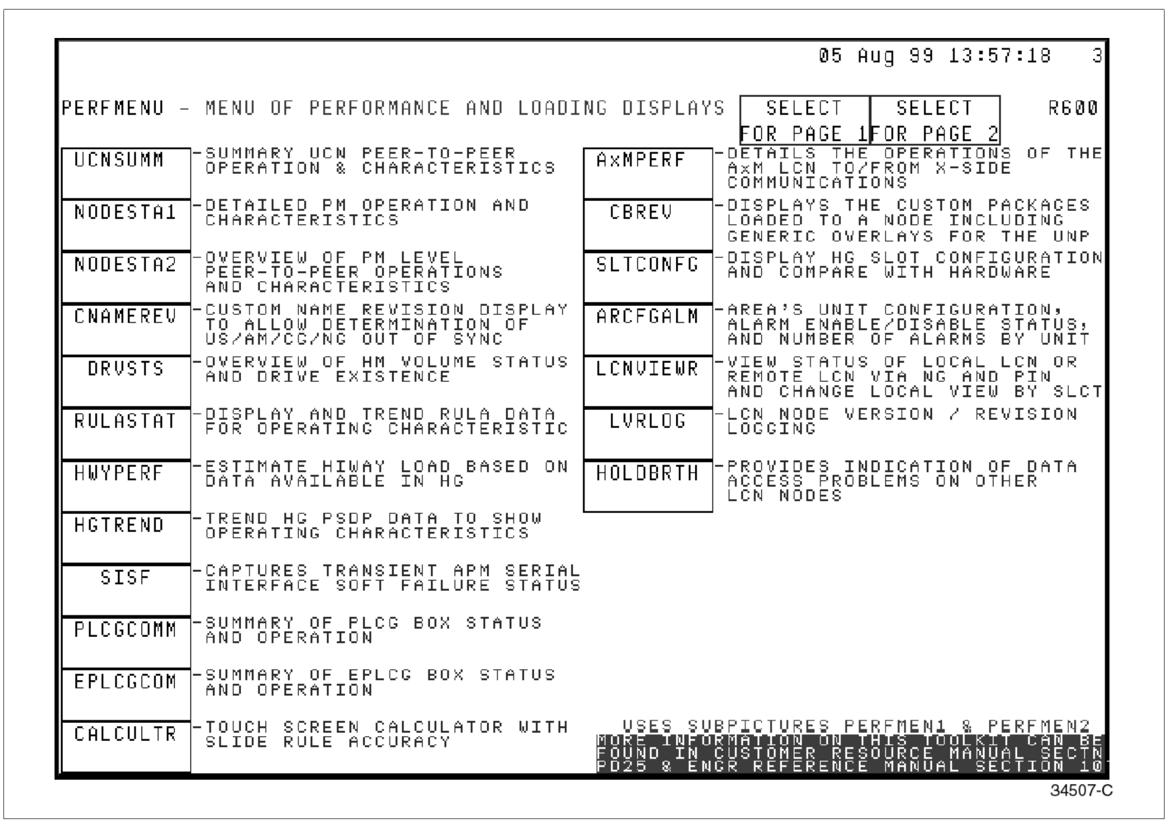

Figure 6 - Menu of Performance and Loading Displays—Page 2

Table 9 - Current TLK1 Versions

Tool Kit Display	Running Version of System Software									
	R400	R401	R410	R420	R430	R431	R500	R510	R520	R600
PERFMENU	R400	R400	R411	R420	R430	R431	R500	R510	R520	R600
NODEPERF	NON E	NONE	R410	R411	R411	R411	R500	R500	R500	R500
DATACHNG	R232	R400	R411	R411	R411	R411	R411	R411	R411	R411
QUIKTRND	R400	R400	R411	R411	R411	R411	R411	R411	R411	R411
CPUCHKR	R400	R400	R410	R410	R410	R410	R410	R410	R410	R410
HEAPCHKR	R400	R400	R411	R411	R411	R411	R411	R411	R411	R411
HEAPMIN	R400	R400	R411	R411	R411	R411	R411	R411	R411	R411
PARCHKR	R400	R400	R400	R400	R400	R400	R400	R400	R400	R400
AMDETAIL	R400	R400	R410	R410	R410	R410	R410	R410	R410	R410
AMTREND	R400	R400	R322	R420	R420	R420	R500	R500	R500	R500
SLTCONFG	R232	R232	R411	R420	R420	R431	R431	R431	T431	R431
HEAPFRAG	R400	R400	R411	R411	R411	R411	R411	R411	R411	R411
CHKPTIME	R300	R300	R411	R411	R411	R411	R411	R411	R411	R411
CLOKMODE	R400	R400	R400	R400	DELETED					
CLOKSYNC	R400	R400	R400	R400	DELETED					
CLOKTRAN	R400	R400	R400	R400	DELETED					
CLOKCABL	R320	R320	R320	R320	DELETED					
UCNCOMM	R321	R321	R411	R411	R411	R411	R500	R500	R500	R500
UCNEVENT	R321	R321	R411	R411	R411	R411	R500	R500	R500	R500
NIMTREND	R322	R322	R322	R322	R322	R322	R500	R500	R500	R500
NGDETAIL	R400	R400	R410	R410	R410	R410	R500	R500	R500	R510
NGTREND	R400	R400	R410	R410	R410	R410	R500	R500	R500	R500
HMDETAIL	R400	R400	R410	R410	R410	R410	R500	R500	R500	R500
HMTREND	R400	R400	R410	R410	R410	R410	R500	R500	R500	R500
HGTREND			R322	R322	R322	R322	R500	R500	R500	R500
UCNSUMM			R411	R411	R411	R411	R411	R411	R411	R411
NODESTA1			R411	R411	R411	R411	R411	R411	R411	R411
NODESTA2			R411	R411	R411	R411	R411	R411	R411	R411
CNAMEREV			R411	R411	R411	R411	R411	R411	R411	R411
DRVSTS			R411	R420	R420	R420	R420	R510	R510	R510
CBREV			R410	R410	R410	R410	R410	R410	R410	R410
IOPMDATA				R322	R322	R322	R322	R322	R322	R322
RULASTAT				R411	R411	R411	R411	R411	R411	R411
HWYPERF				R323	R323	R323	R323	R323	R323	R323
SISF				R420	R420	R420	R420	R420	R420	R420
UCNVERSN				R420	R420	R420	R420	R420	R420	R420
PLCGCOMM				R430	R430	R430	R430	R431	R510	R510
CALCULTR					R420	R431	R431	R431	R431	R431
CLOKSTAT					R430	R431	R431	R431	R431	R431
AXMPERF						R431	R500	R500	R500	R500
HISGRPS							R430	R430	R430	R531
AMDIAGNS								R510	R510	R510
ARCFGALM								R510	R510	R510
EPLCGCOM									R510	R510
LCNVIEWR									R520	R520
LVRLOG										R430
HOLDBRTH										R530

Node Performance Display

Figures 7 – 13 show the NODEPERF displays for several types of nodes:

Figure 7—AM

Figure 8—HM

Figure 9—HG

Figure 10—NIM

Figure 11—US

Figure 12—NG

Figure 13—CG

The NODEPERF display provides a view into the PSDP data on each node of the system. Not all available information is shown, just that data of greatest interest when attempting to determine the operational characteristics of that node.

A node can be selected by touching the SELECT FOR NODE target and entering the node to be monitored. (NOTE: The first six characters of the PSDP point name are already typed-in, so the display can be used without an Engineering keyboard.)

The upper part of the display contains parameters and information unique to each node type. The bottom part the display contains common information for all nodes.

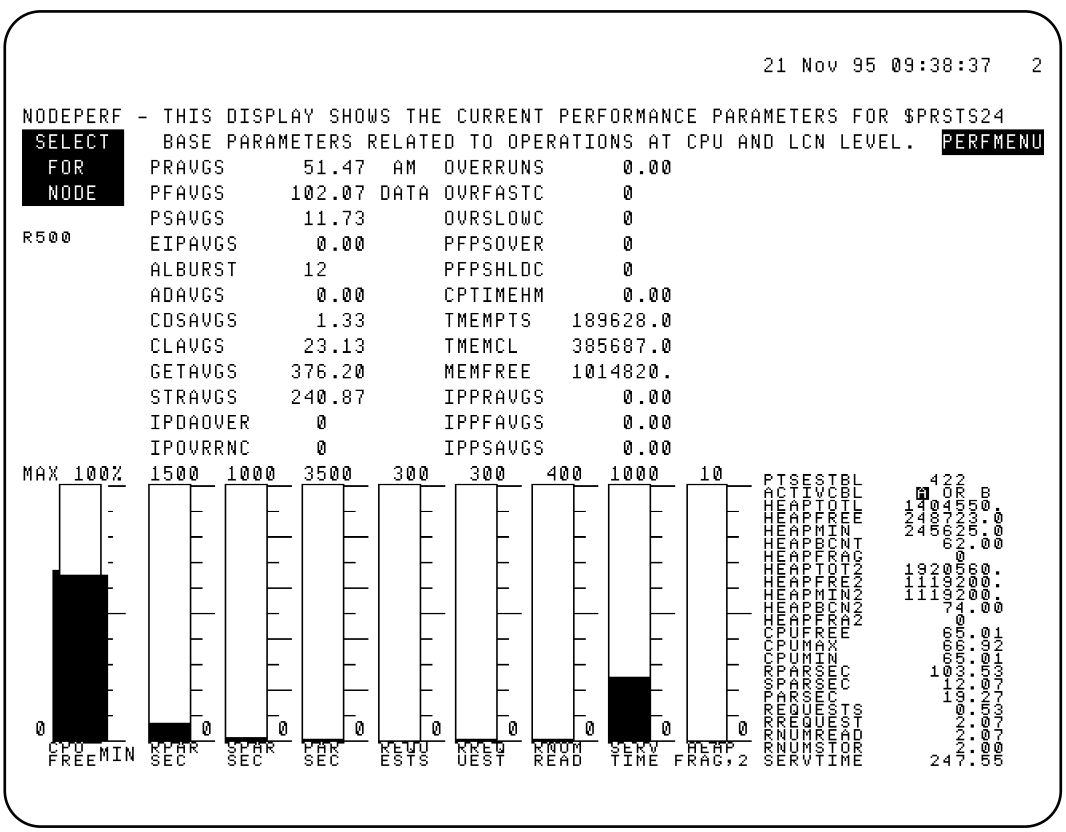

34508

Figure 7 - AM NODEPERF Display

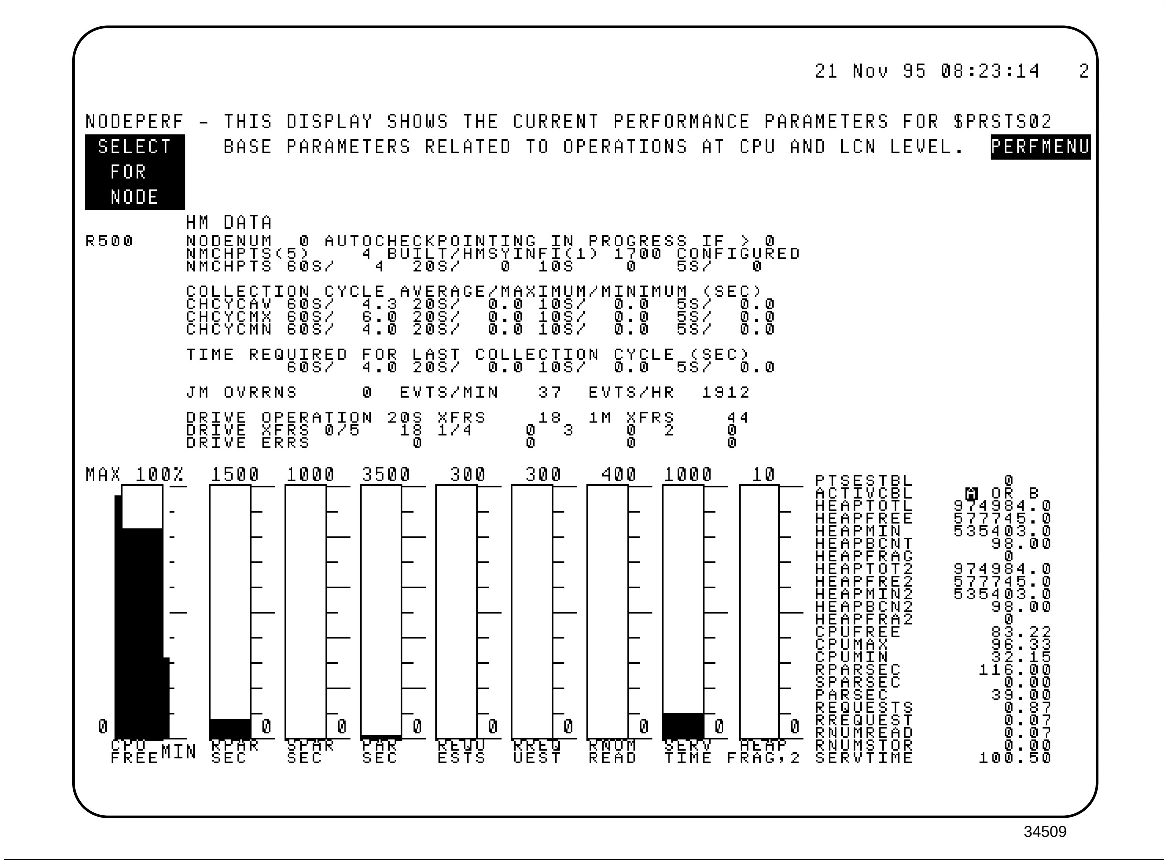

Figure 8 - HM NODEPERF Display

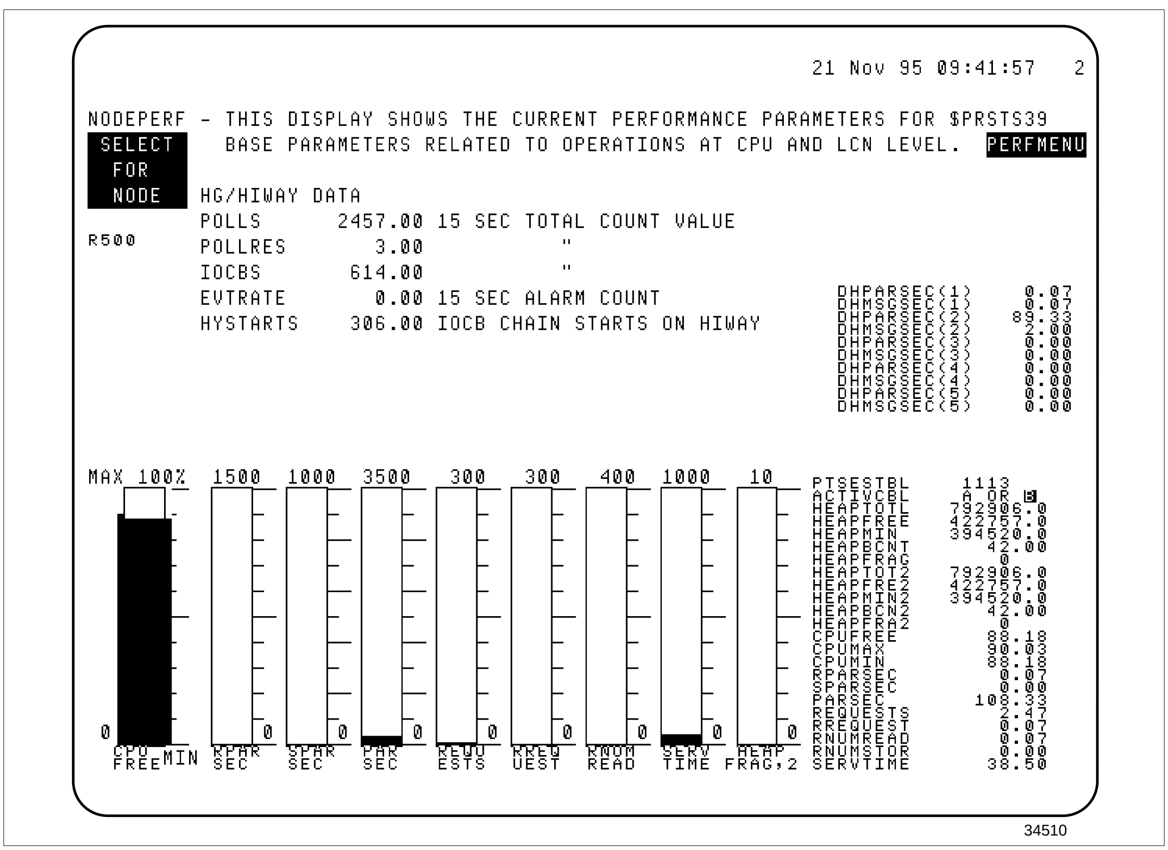

Figure 9 - HG NODEPERF Display

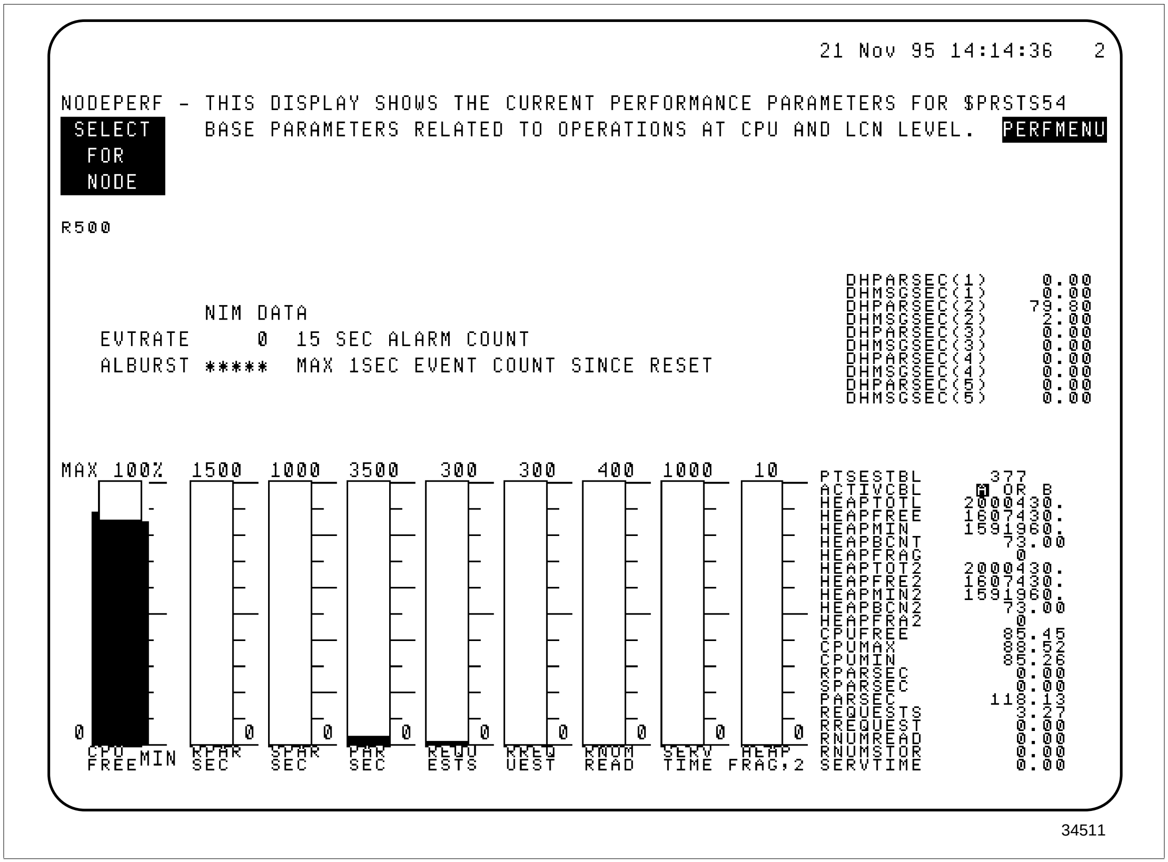

Figure 10 - NIM NODEPERF Display

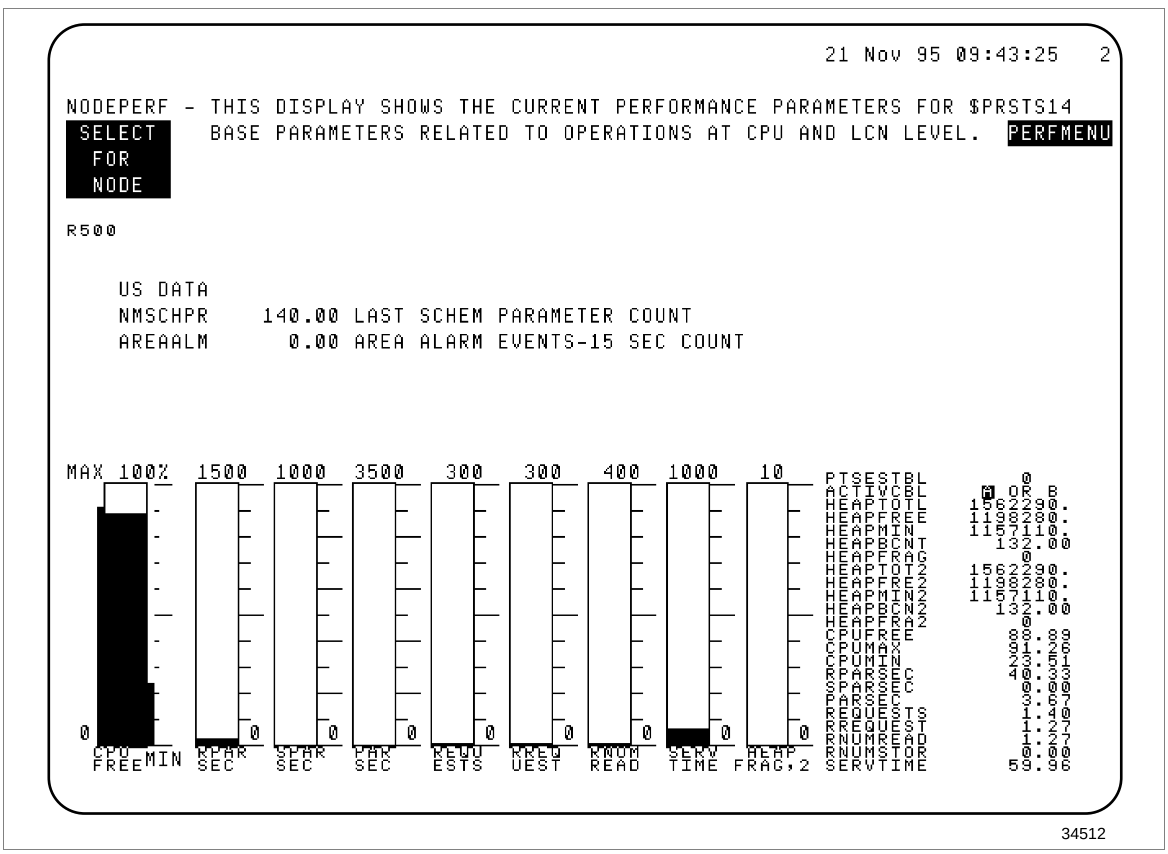

Figure 11 - US NODEPERF Display

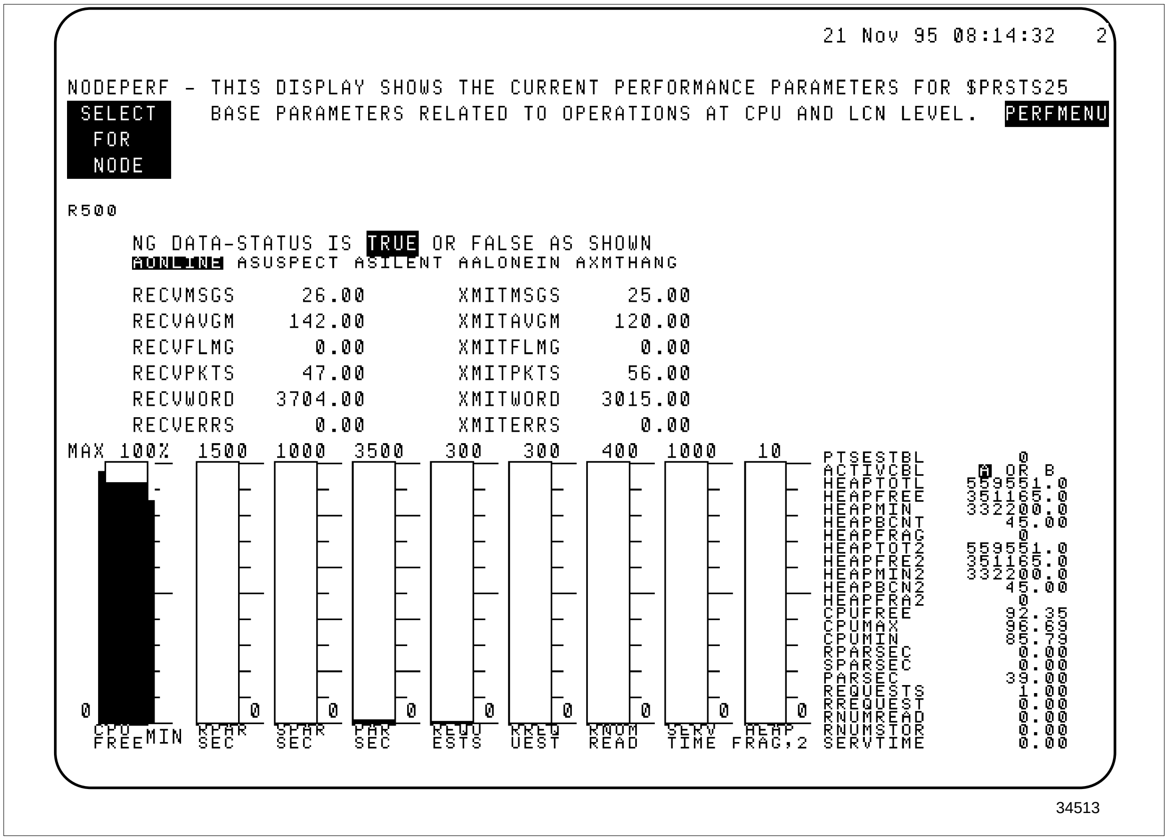

Figure 12 - NG NODEPERF Display

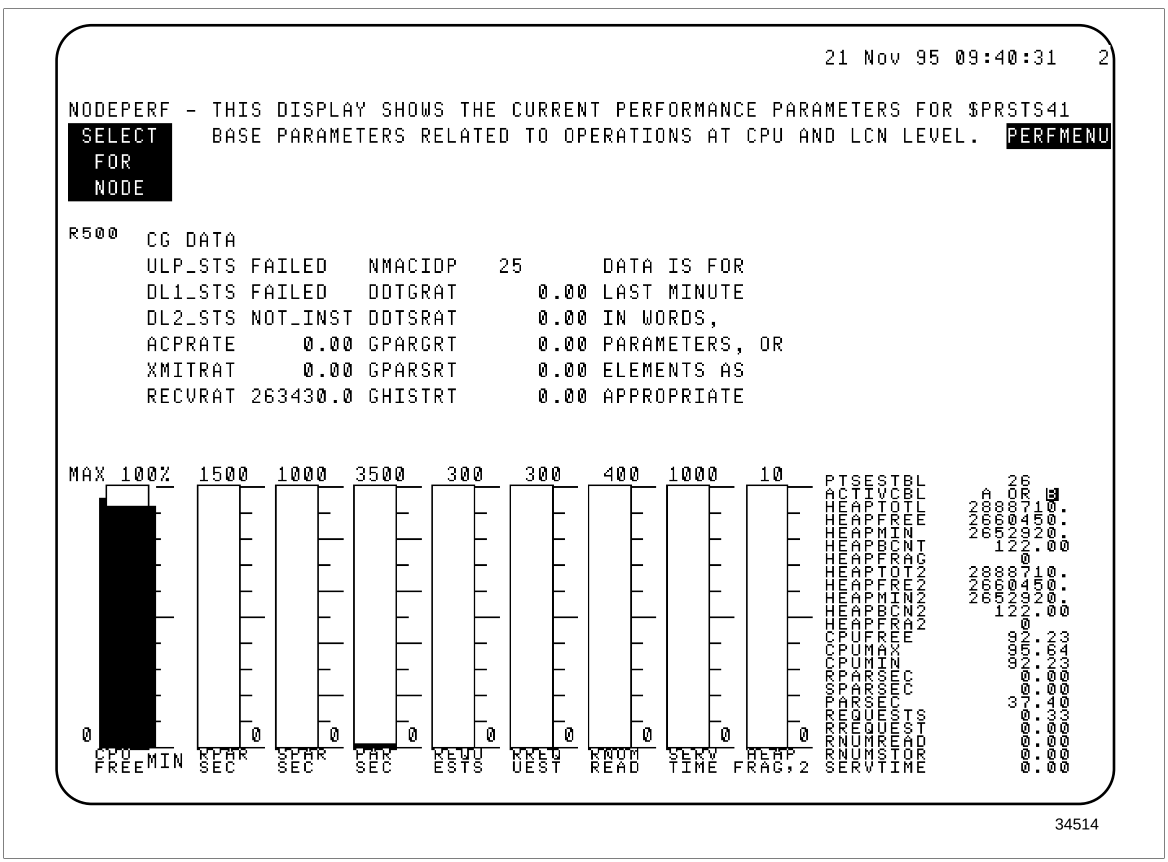

Figure 13 - CG NODEPERF Display

Quick Trend Display

The QUIKTRND display allows you to trend up to four variables of PSDP data or other parameters.

This trend is particularly useful for quick trending of suspicious nodes. Two associated displays, QUIKHLP1 and QUIKHLP2, provide convenient help information.

The data entered in the QUIKTRND display is saved in memory so that recall and restart of the trend is simplified. The R400 version of this display provides a scrolling capability.

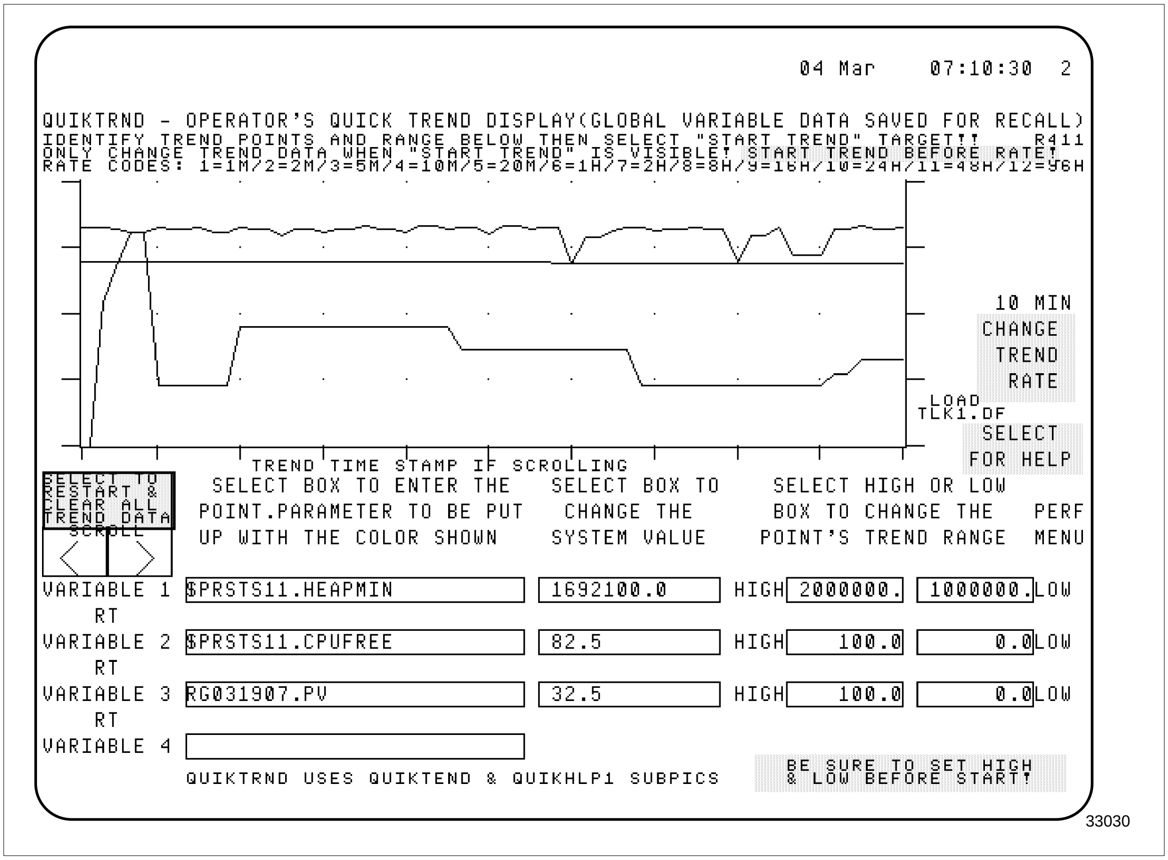

Figure 14 - QUIKTRND Display

Each of the four variables have individually configurable ranges. The range must be defined before starting the trend. Each display has a configurable rate entered through a rate code.

Data Change Display

The DATACHNG display provides generic access to the “point.parameter” data in the TPS Network system.

Up to 10 values can be monitored and/or changed and the data held as global data until the item is reset or the US is reloaded. This allows you to set up a list of data to monitor, then return later without having to retype the desired information.

When the item is selected for the point.parameter entry, the port is set up with the characters in the position to allow entry of the PSDP and indexed parameter without using the engineering keyboard.

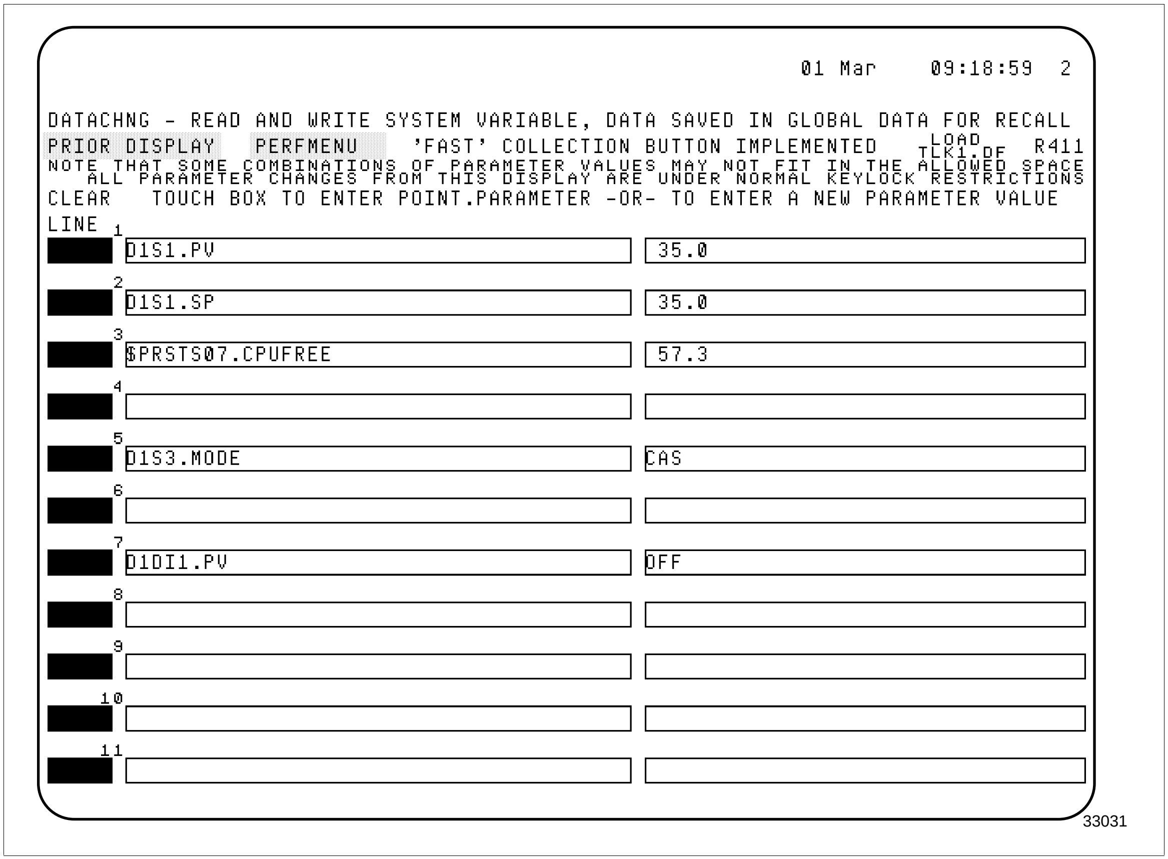

Figure 15 - DATACHNG Display

CPU Checker Display

The CPUCHKR display shows the percentage of free time available in the CPU of every node on the system simultaneously.

The CPUFREE percentage is very helpful in determining a specific node's load during the last 15 seconds. The percentage indicates how close the node is to the processing limit of the processor board that is installed in that node.

```
                                                           10 Nov    10:08:15  4

CPUCHKR - CHECKS EACH NODE'S CPUFREE (%) AGAINST AN ENTERED CHECK VALUE    R410
          8 SECOND DISPLAY UPDATE, VALUE CALCUALTED EVERY 15 SECONDS
ND 01  54.80   ND 21  93.39   ND 41  55.79   ND 61  80.25   ND 81
ND 02  64.38   ND 22  10.94   ND 42  93.86   ND 62  41.88   ND 82
ND 03  34.33   ND 23  93.18   ND 43  64.81   ND 63  81.31   ND 83
ND 04  37.73   ND 24  75.72   ND 44  93.03   ND 64  29.69   ND 84
ND 05  45.84   ND 25  97.86   ND 45  83.69   ND 65  64.56   ND 85  36.79
ND 06  39.76   ND 26  99.47   ND 46  30.73   ND 66  82.99   ND 86  80.48
ND 07  30.64   ND 27  97.85   ND 47  81.65   ND 67          ND 87
ND 08  33.13   ND 28  98.22   ND 48  21.38   ND 68          ND 88  82.65
ND 09  57.95   ND 29  51.00   ND 49  70.00   ND 69          ND 89
ND 10  75.02   ND 30  51.27   ND 50  55.37   ND 70  96.77   ND 90  85.39
ND 11  72.00   ND 31  87.47   ND 51  68.84   ND 71  98.58   ND 91  85.55
ND 12  74.35   ND 32  94.94   ND 52  35.87   ND 72  98.67   ND 92  83.59
ND 13  60.51   ND 33  40.95   ND 53  54.91   ND 73          ND 93  78.02
ND 14  59.04   ND 34  39.95   ND 54  48.23   ND 74          ND 94  73.87
ND 15  27.55   ND 35  53.10   ND 55  65.05   ND 75  80.90   ND 95  77.78
ND 16  38.48   ND 36  41.20   ND 56  25.41   ND 76          ND 96  74.12
ND 17  76.51   ND 37  51.43   ND 57  82.91   ND 77
ND 18  77.39   ND 38  66.29   ND 58  27.66   ND 78
ND 19  78.02   ND 39   1.94   ND 59  100.0   ND 79                 PARCHKR
ND 20  12.34   ND 40  27.70   ND 60  54.01   ND 80                 DISPLAY
                                                                   HEAPCHKR
                      25% TEST VALUE                               DISPLAY
$PRSTS16     CHECKER     SELECT CHECKER TO ENTER LOW CPU % LIMIT   PFRFMFNII
   ON        TARGET   VALUE.  NODE'S VALUE TURNS RED WHEN <= LIMIT. DISPLAY
```

33032

Figure 16 - CPUCHKR Display

Heap Checker Display

The HEAPCHKR display provides a convenient display to observe the current values of heap memory available in each node.

A target can be selected to allow the user to enter a value, and if any node's available heap memory should go equal to or below that value it would be displayed in red.

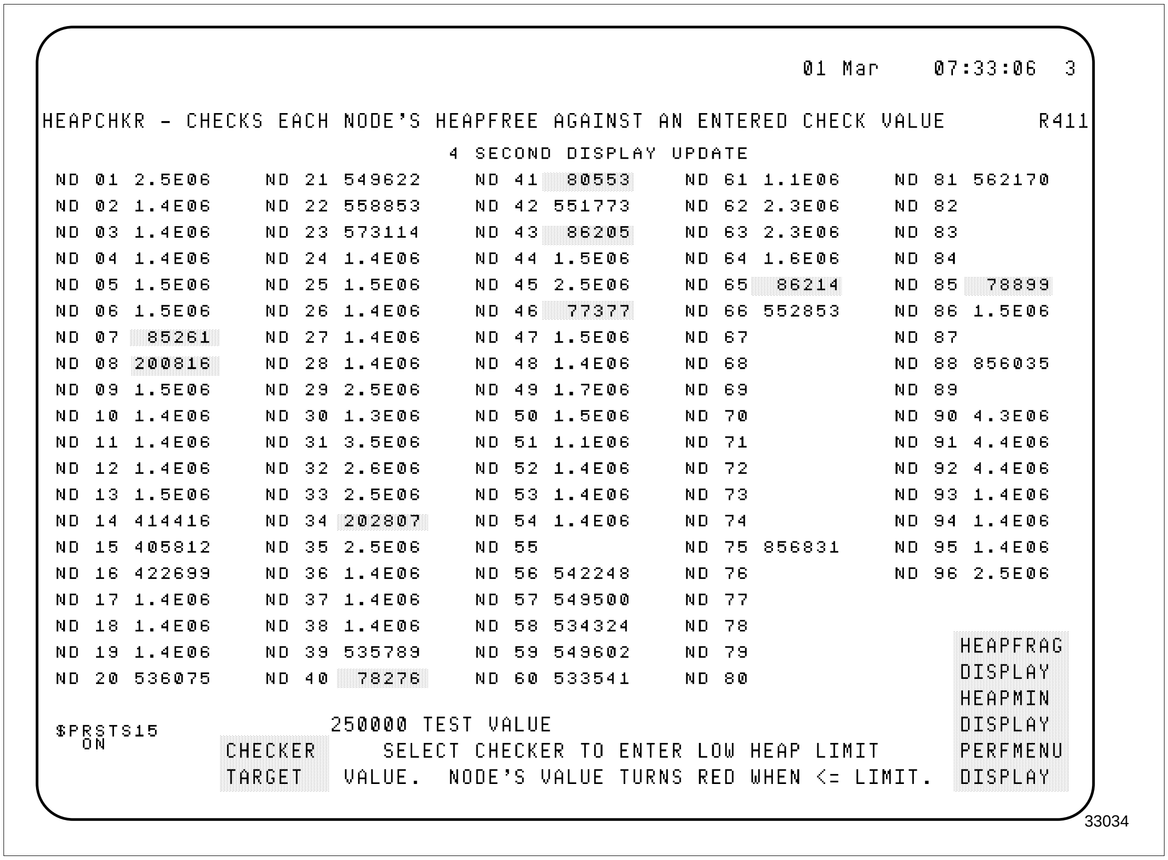

Figure 17 - HEAPCHKR Display

Heap Minimum Display

The HEAPMIN display shows the lowest amount of heap memory that is available in each node of the system since that node was loaded.

A target can be selected to allow the user to enter a value, and if any node’s available heap memory should go equal to or below that value it would be displayed in red.

Both POOL1 and POOL2 heap values are shown for each node:

- POOL1 values—full intensity green
- POOL2 values—half intensity green

Only US/GUSs in an engineering function and AMs use both pools of heap. For all other node types, there is only one heap pool and both numbers are identical.

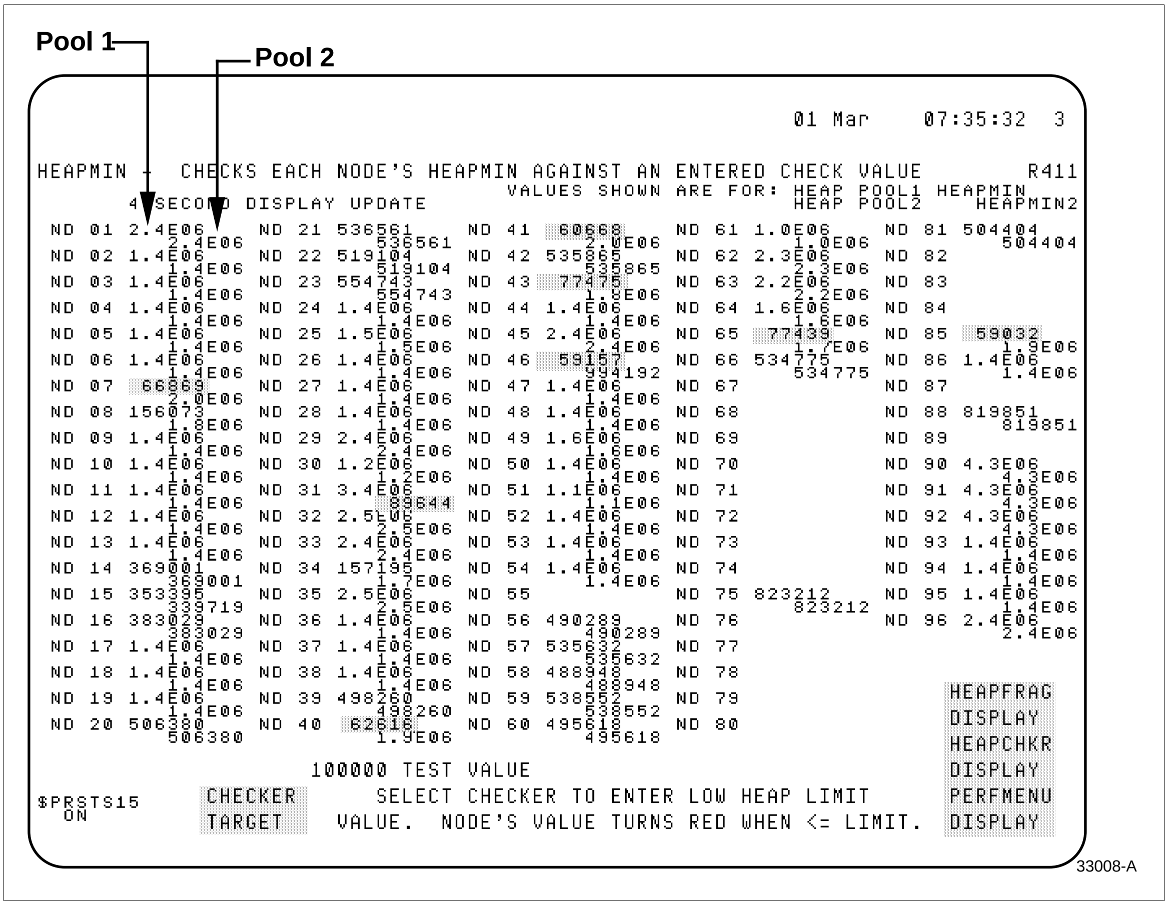

Figure 18 - HEAPMIN Display

Heap Fragmentation Display

The HEAPFRAG display shows heap fragmentation. HEAPFRAG values of 1 to 3 are expected in US nodes running engineering functions, such as DEB, CL, and PE.

A target can be selected to allow the user to enter a value and if any node's available heap fragmentation should go equal to or above that value it would be displayed in red.

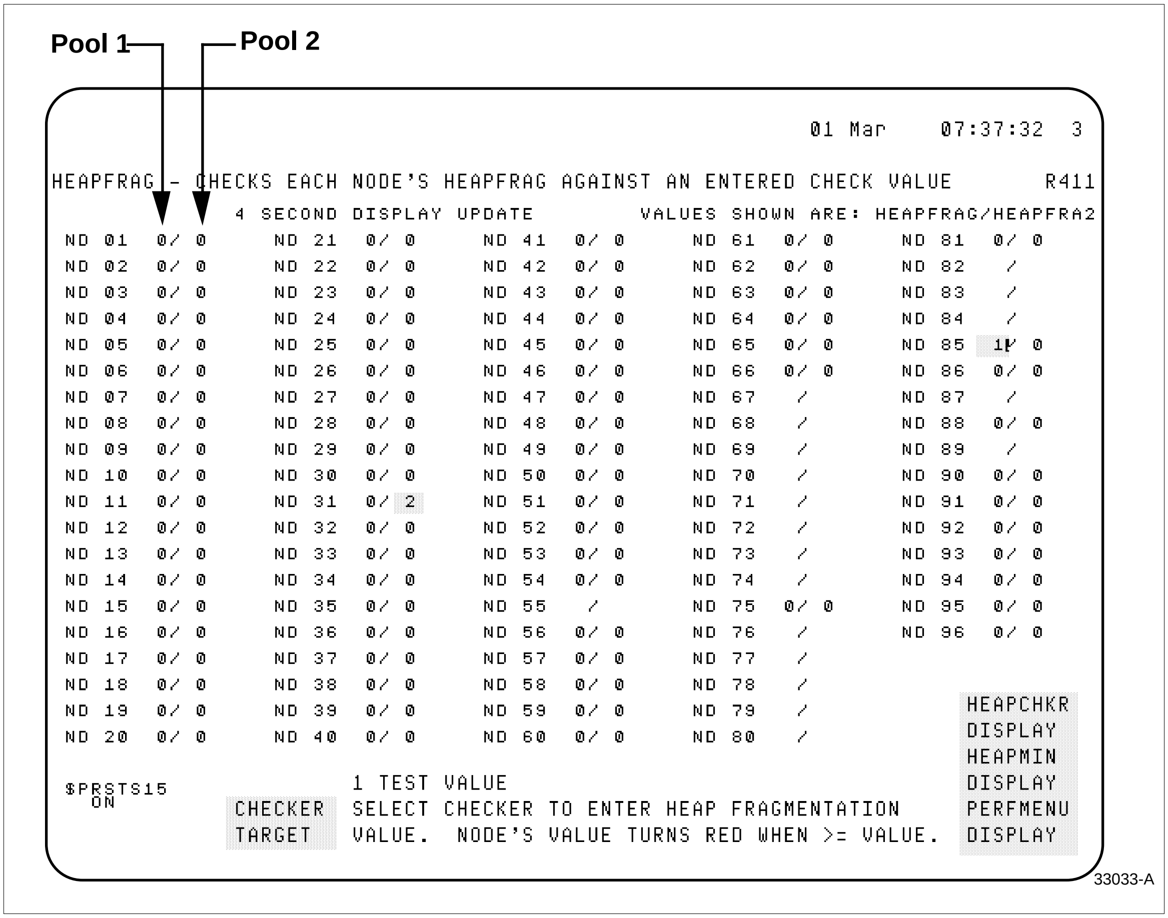

Figure 19 - HEAPFRAG Display

Parameter Checker Display

The PARCHKR display shows the number of parameters per second being stored to or read from each node by other LCN nodes on the system.

Each node on the system is shown simultaneously.

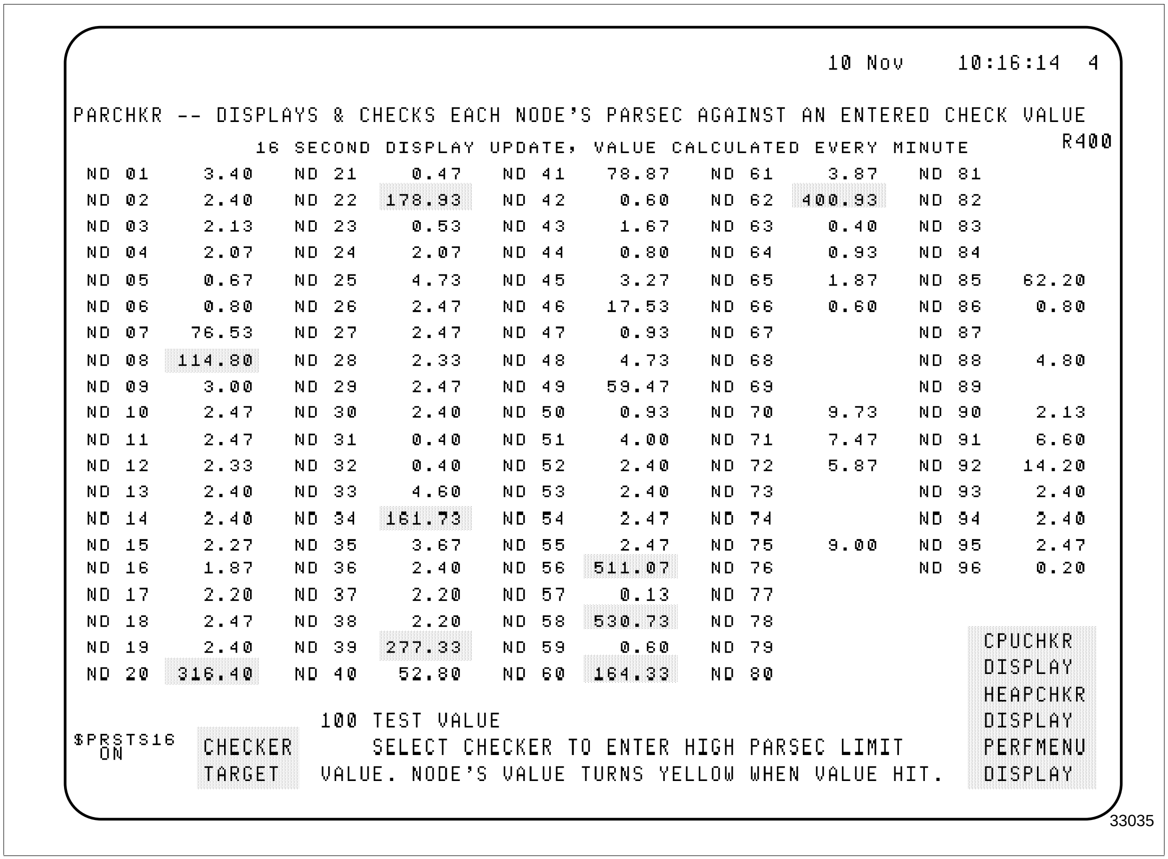

10 Nov 10:16:14 4

PARCHKR -- DISPLAYS & CHECKS EACH NODE'S PARSEC AGAINST AN ENTERED CHECK VALUE

16 SECOND DISPLAY UPDATE, VALUE CALCULATED EVERY MINUTE R400

ND 01	3.40	ND 21	0.47	ND 41	78.87	ND 61	3.87	ND 81	
ND 02	2.40	ND 22	178.93	ND 42	0.60	ND 62	400.93	ND 82	
ND 03	2.13	ND 23	0.53	ND 43	1.67	ND 63	0.40	ND 83	
ND 04	2.07	ND 24	2.07	ND 44	0.80	ND 64	0.93	ND 84	
ND 05	0.67	ND 25	4.73	ND 45	3.27	ND 65	1.87	ND 85	62.20
ND 06	0.80	ND 26	2.47	ND 46	17.53	ND 66	0.60	ND 86	0.80
ND 07	76.53	ND 27	2.47	ND 47	0.93	ND 67		ND 87	
ND 08	114.80	ND 28	2.33	ND 48	4.73	ND 68		ND 88	4.80
ND 09	3.00	ND 29	2.47	ND 49	59.47	ND 69		ND 89	
ND 10	2.47	ND 30	2.40	ND 50	0.93	ND 70	9.73	ND 90	2.13
ND 11	2.47	ND 31	0.40	ND 51	4.00	ND 71	7.47	ND 91	6.60
ND 12	2.33	ND 32	0.40	ND 52	2.40	ND 72	5.87	ND 92	14.20
ND 13	2.40	ND 33	4.60	ND 53	2.40	ND 73		ND 93	2.40
ND 14	2.40	ND 34	161.73	ND 54	2.47	ND 74		ND 94	2.40
ND 15	2.27	ND 35	3.67	ND 55	2.47	ND 75	9.00	ND 95	2.47
ND 16	1.87	ND 36	2.40	ND 56	511.07	ND 76		ND 96	0.20
ND 17	2.20	ND 37	2.20	ND 57	0.13	ND 77			
ND 18	2.47	ND 38	2.20	ND 58	530.73	ND 78			
ND 19	2.40	ND 39	277.33	ND 59	0.60	ND 79			
ND 20	316.40	ND 40	52.80	ND 60	164.33	ND 80			

CPUCHKR DISPLAY
HEAPCHKR DISPLAY
PERFMENU DISPLAY

100 TEST VALUE

$PRSTS16 ON

CHECKER TARGET

SELECT CHECKER TO ENTER HIGH PARSEC LIMIT VALUE. NODE'S VALUE TURNS YELLOW WHEN VALUE HIT.

33035

Figure 20 - PARCHKR Display

AM Detail Display

The AMDETAIL display allows you to view a large number of the significant AM PSDP values.

The AM PSDP values are shown as snapshot values, previous hour values, and current hour values. Some of the standard values are also given for a total view of the AM loading.

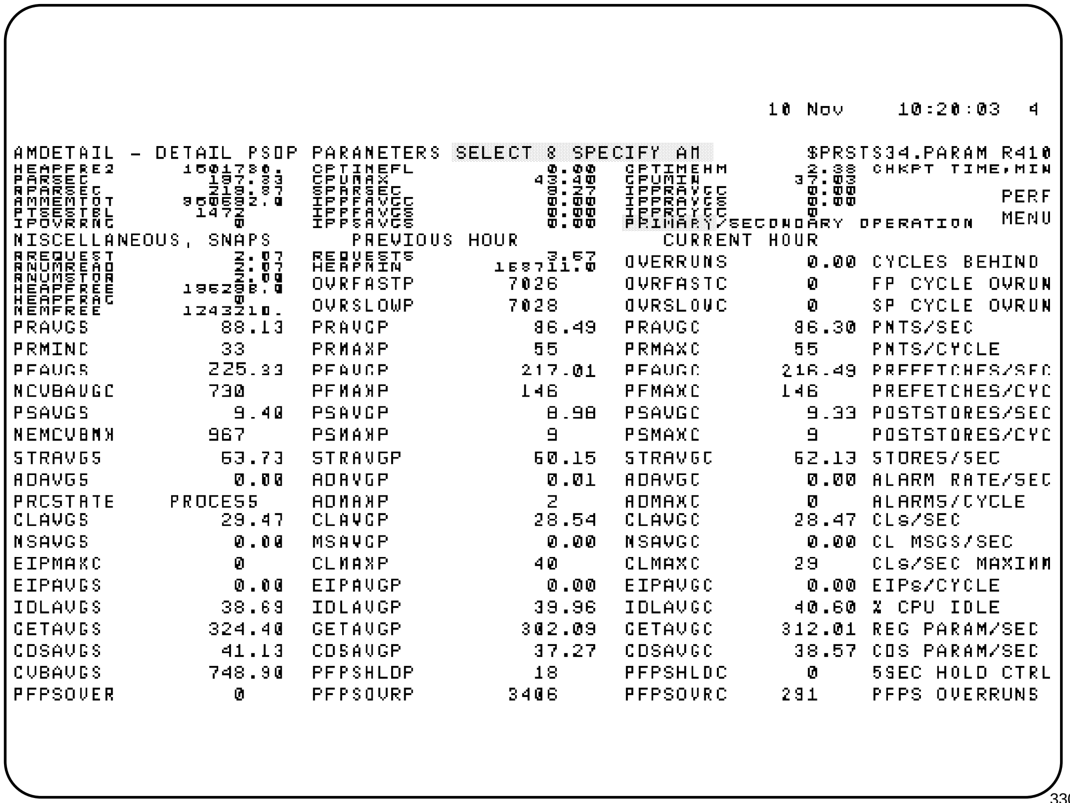

Figure 21 - AMDETAIL Display

AM Trend Display

The AMTREND display provides a trend of 16 of the most significant indicators of AM loading and performance.

A SELECT & SPECIFY AM target is at the top of the screen. When this target is selected, it opens the port for the AM's node number to be entered. A "Start Trend" starts the trending with a default time base of 24 hours.

After the trend is started, the CHANGE RATE target can be selected and the rate can be changed.

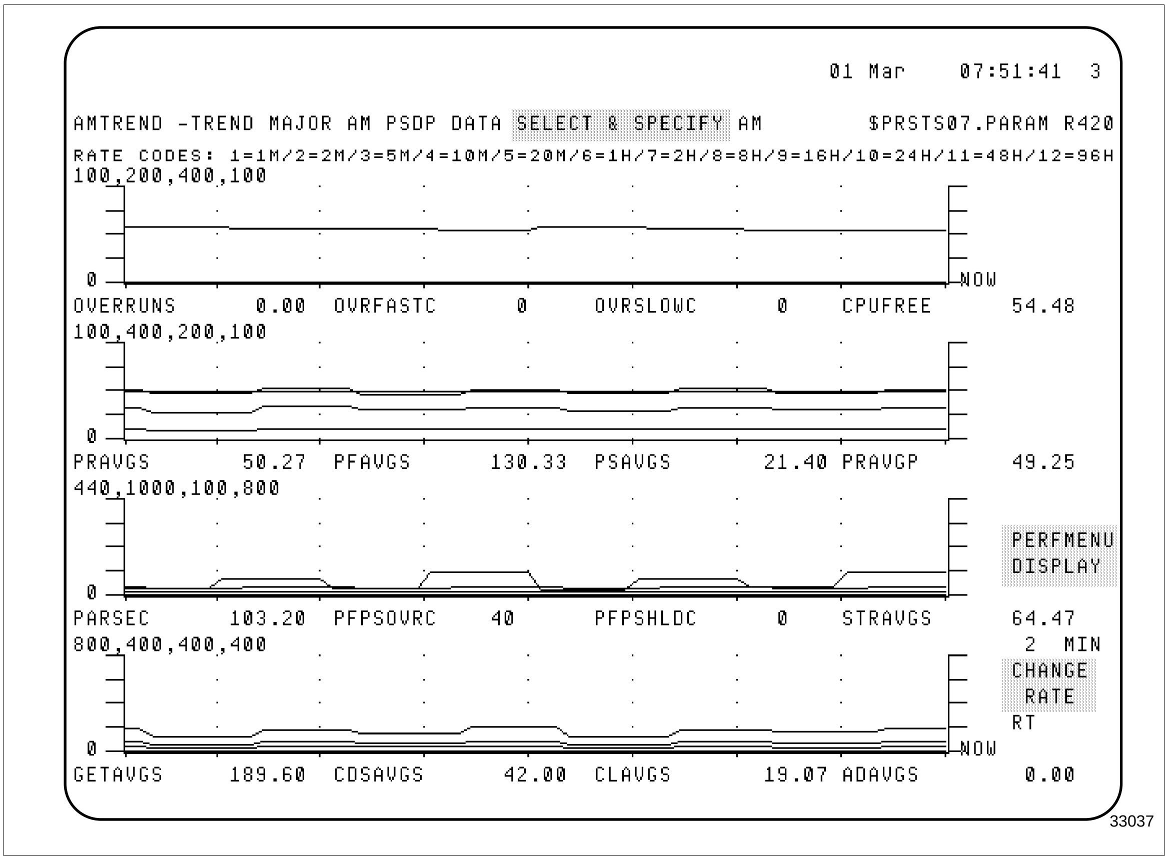

Figure 22 - AMTREND Display

AM Diagnostic Display

Parameters have been added to the AM in R510 to allow better determination of AM operations at the CL, fast, and slow processor level. Some of these parameters are saved and made available to the SMCC for display if the AM crashes. The AMDIAGNS display provides on line access to this data.

The display has a ’FST’ collection group, so that the FAST button (or FAST target) can be used to obtain rapid display update. Additionally, an ON/OFF target is provided to allow additional AM diagnostic data to be displayed.

CAUTION—When the DIAGSTAT switch is ON, the display imposes a significant load on the AM.

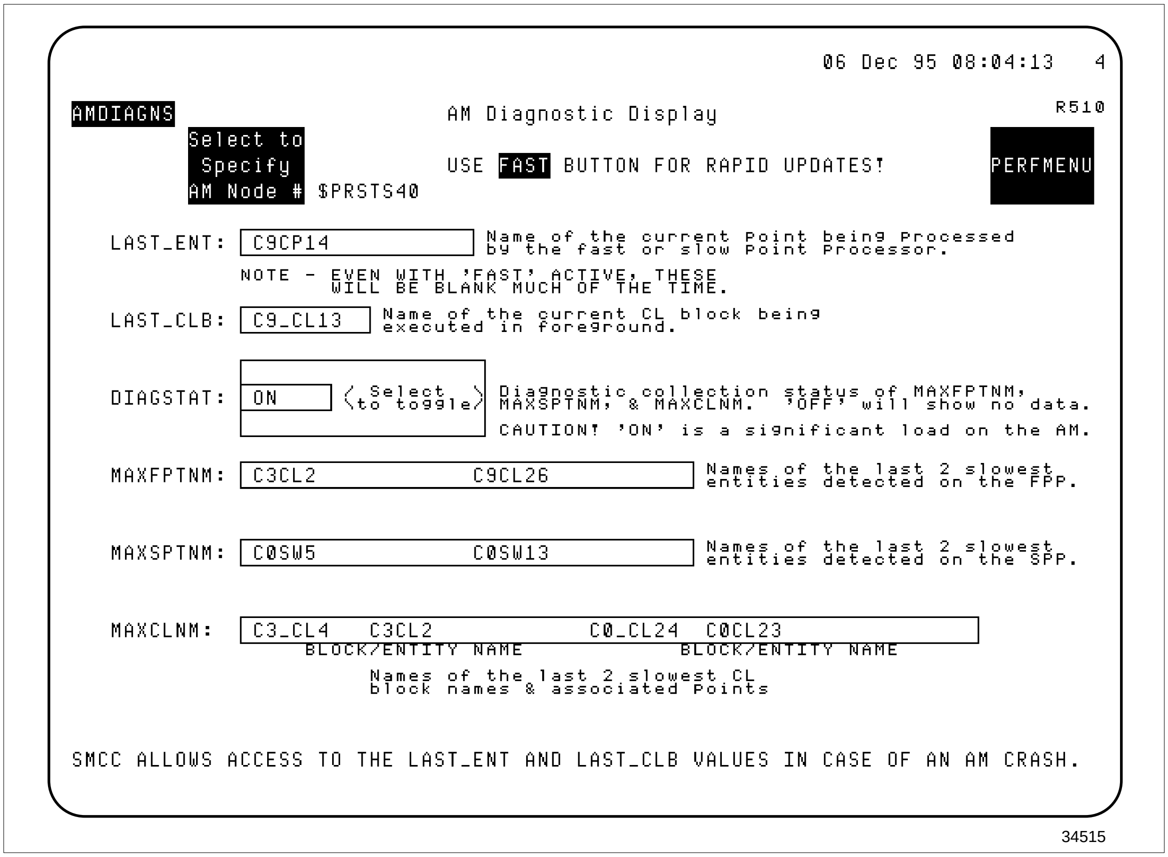

Figure 23 - AMDIAGNS Display

HG Trend Display

The HGTREND display provides trend data for two HGs. Two valid HG nodes must be specified (use the same number for both, if necessary).

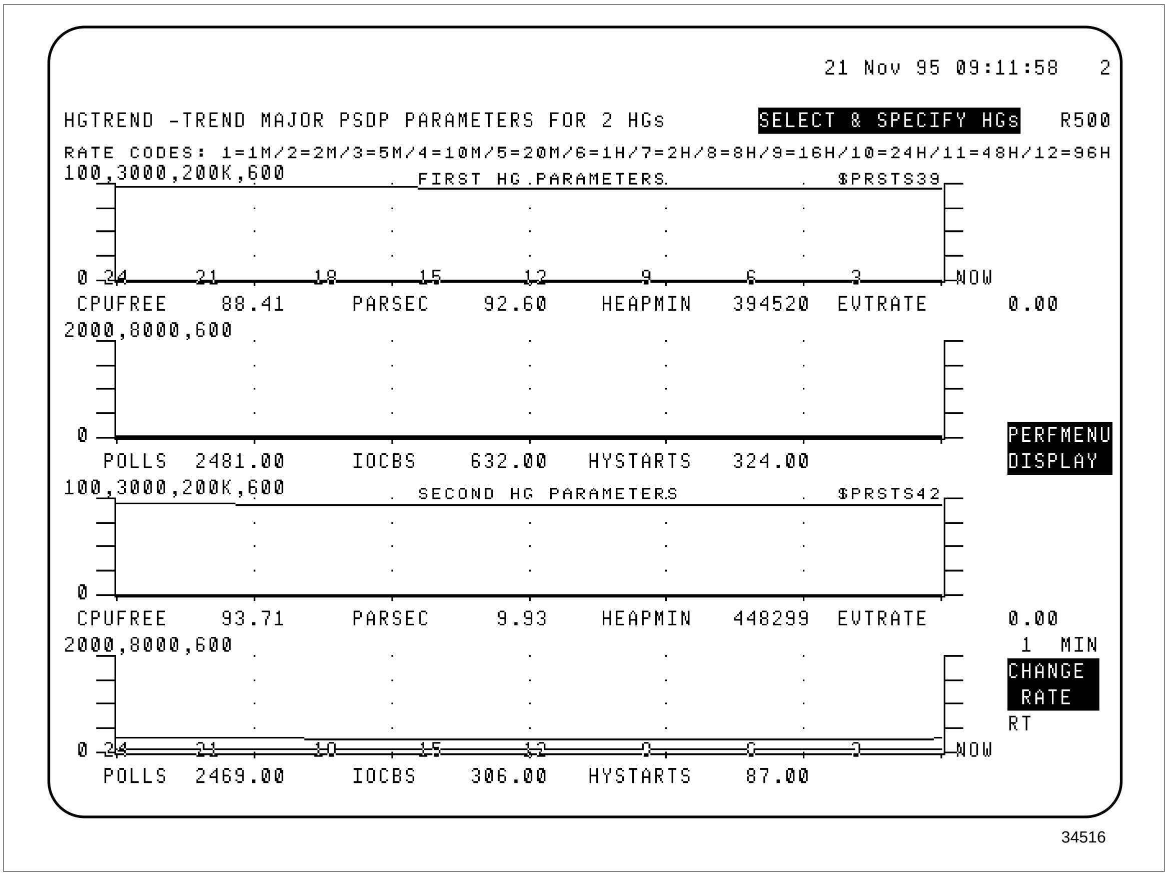

Figure 24 - HGTREND Display

NG Detail Display

The NGDETAIL display provides all the detail available on the Network Gateway. A small trend grid is provided for common items.

A second page shows Plant Information Network (PIN) driver details.

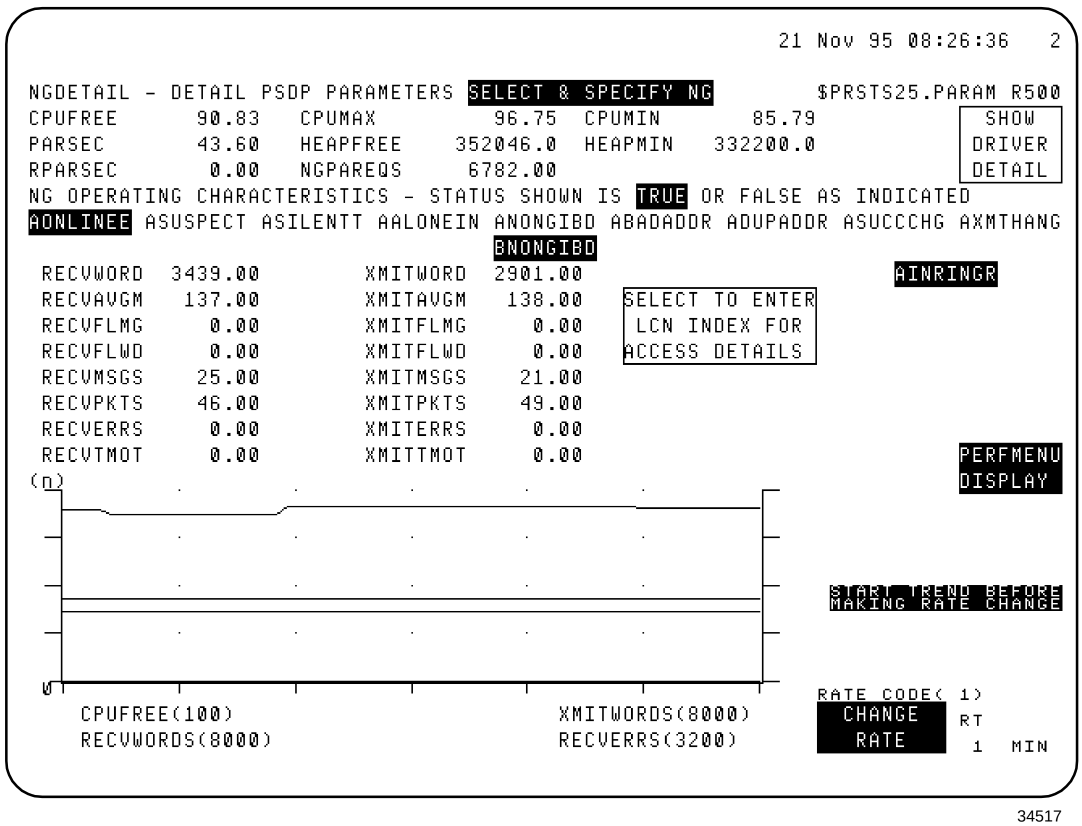

Figure 25 - NGDETAIL Display—NG Node Selected

21 Nov 95 08:29:32 2

NGDETAIL - DETAIL PSDP PARAMETERS SELECT & SPECIFY NG $PRSTS25.PARAM R500

PROVIDES DETAILS ON NG DRIVER OPERATION - ACBLSTAT/BCBLSTAT(1-48)

SHOW MAIN PAGE

	CABLE A	CABLE B		CABLE A	CABLE B
TOKENS PASSED	0	0	REMOD EXCEEDED	0	0
TOKENS HEARD	0	0	FRAGMENT EXCEED	0	0
NO SUCCESSOR	0	0	NO Rx BUFFER	0	0
WHO FOLLOWS	0	0	Rx INACTIVE NOD	0	0
TOKEN PASS FAIL	0	0	Rx INACTIVE SAP	0	0
NOISE BURST	0	0	LLC SPARE 3	0	0
CHECKSUM ERROR	0	0	LLC SPARE 2	0	0
e-BIT ERROR	0	0	LLC SPARE 1	0	0
FRAME FRAGMENT	0	0	Tx 0 PACKET	0	0
Rx FRAME LONG	0	0	Tx 1 PACKET	36	0
NO Rx RESOURCES	0	0	Tx 2 PACKET	0	0
Rx OVERRUN	0	0	Tx 3 PACKET	9	0
DUPLICATE RWR	0	0	Tx NO BUFFER	0	0
NULL RWR	0	0	Tx SPARE 2	80	0
NO RESPONSE RWR	0	0	Tx SPARE 1	80	0
UNXP RESPNS RWR	0	0	Rx 0 PACKET	0	0
ERR RESPNS RWR	0	0	Rx 1 PACKET	36	0
Tx UNDERRUN	0	0	Rx 2 PACKET	0	0
TBC SPARE 1	0	0	Rx 3 PACKET	5	0
TOKEN ROT TIME	106	0	Rx BAD MSG	0	0
TOKN ROT EXCEED	0	0	Rx SPARE 2	0	0
TOKN FAIL EXCED	0	0	Rx SPARE 1	0	0
NOISE EXCEEDED	0	0	COM SPARE 2	80	0
FCS EXCEEDED	0	0	COM SPARE 1	80	0

34518

\

Figure 26 - NGDETAIL Display—Driver Detail Page

NG Trend Display

The NGTREND display provides trends for major Network Gateway PSDP data. (NOTE: The values in the trend display are running totals and are not reset every few minutes as other trend displays are.)

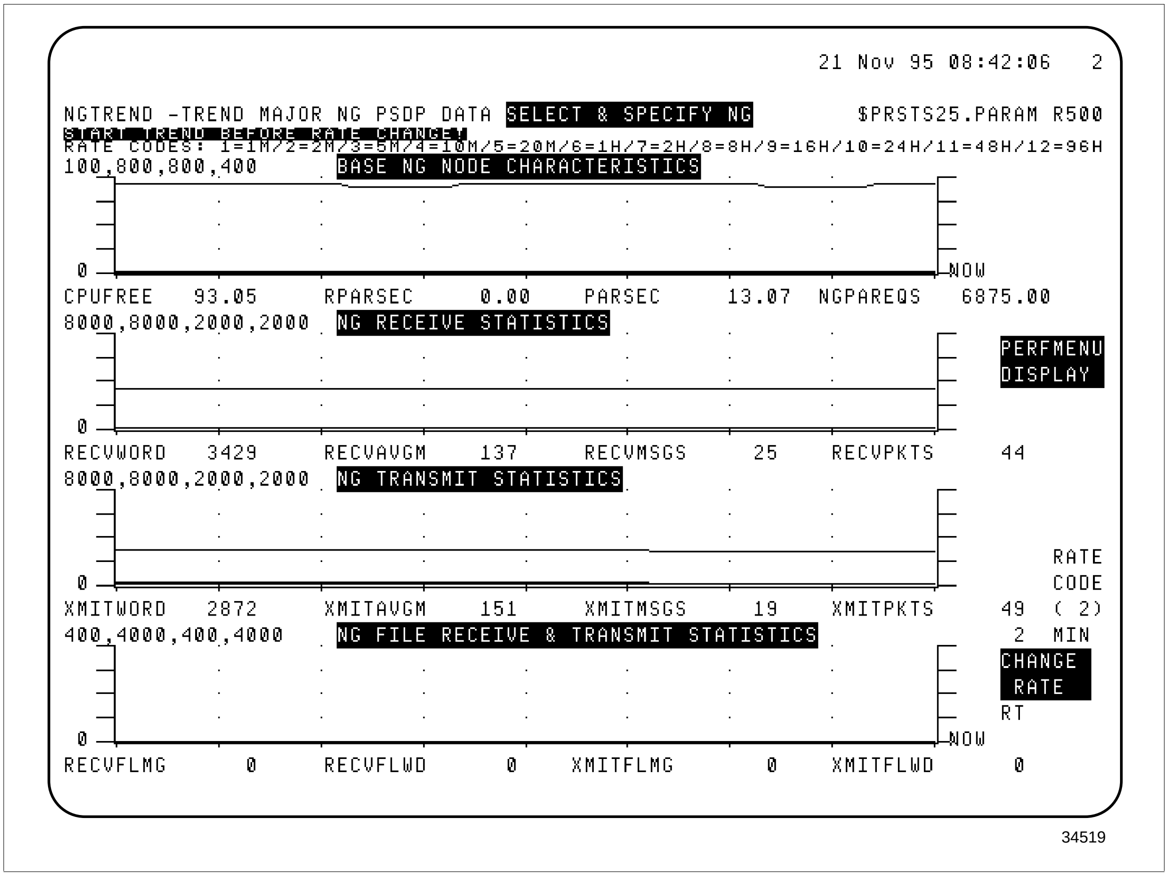

Figure 27 - NGTREND Display

Appendix A describes the parameters seen on these displays.

HM Detail Display

The HMDETAIL display provides detail information on HM operation, including transfers, errors, and a small trend grid for common items.

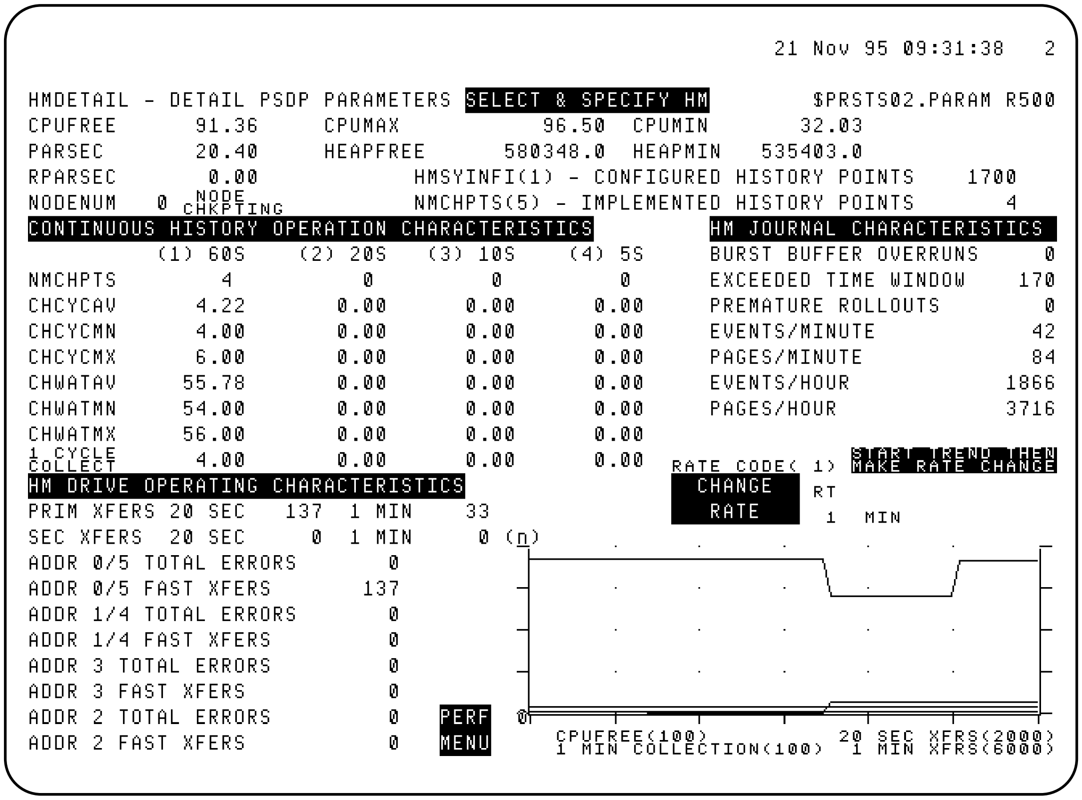

34520

Figure 28 - HMDETAIL Display

HMTREND Display

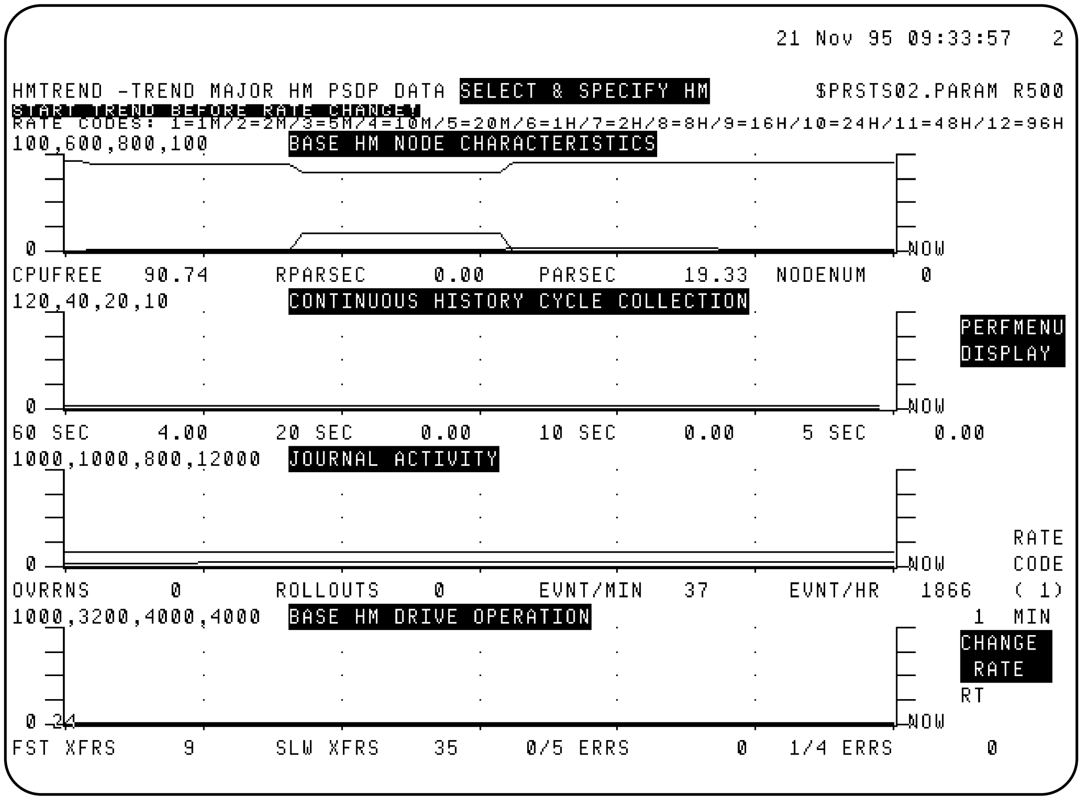

34521

Figure 29 - HMTREND Dis

LCNVIEWR Display

This display provides the view of either the local or remote LCN from any particular node.

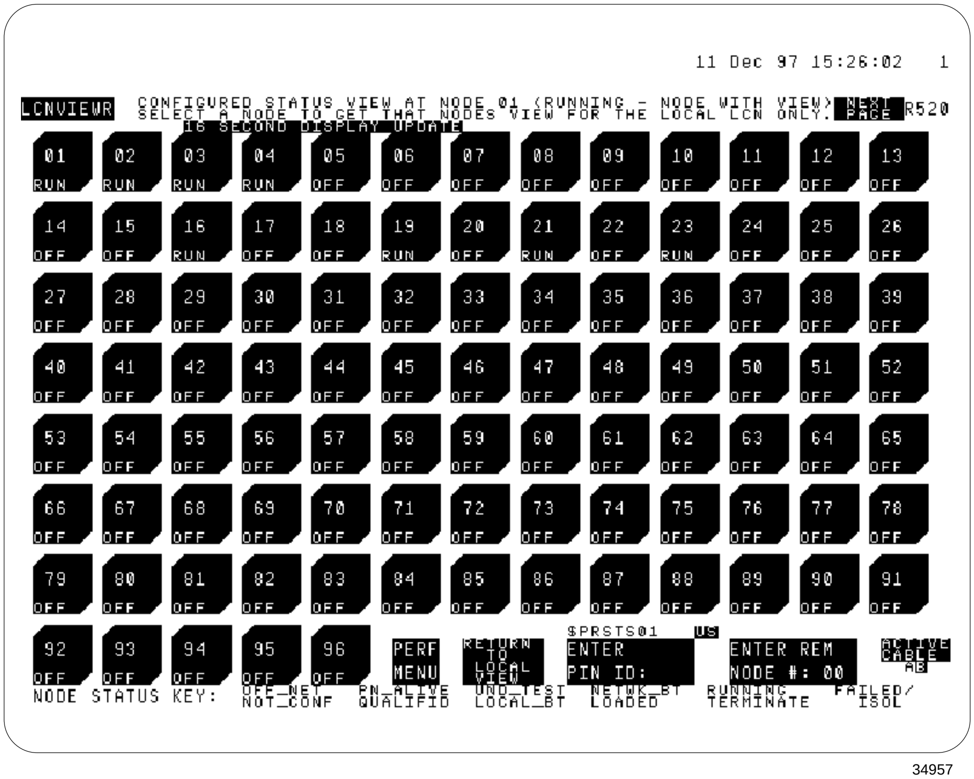

Figure 30 - LCNVIEWR Dis

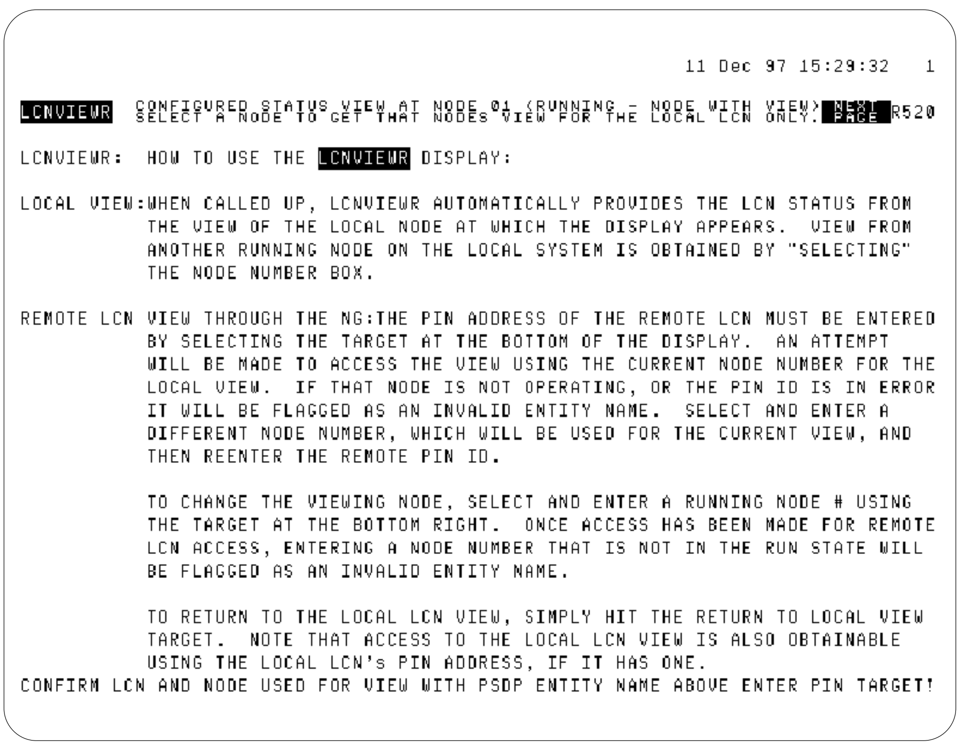
11 Dec 97 15:29:32 1

LCNVIEWR CONFIGURED STATUS VIEW AT NODE 01 (RUNNING - NODE WITH VIEW) NEXT PAGE R520
SELECT A NODE TO GET THAT NODES VIEW FOR THE LOCAL LCN ONLY.

LCNVIEWR: HOW TO USE THE LCNVIEWR DISPLAY:

LOCAL VIEW: WHEN CALLED UP, LCNVIEWR AUTOMATICALLY PROVIDES THE LCN STATUS FROM THE VIEW OF THE LOCAL NODE AT WHICH THE DISPLAY APPEARS. VIEW FROM ANOTHER RUNNING NODE ON THE LOCAL SYSTEM IS OBTAINED BY "SELECTING" THE NODE NUMBER BOX.

REMOTE LCN VIEW THROUGH THE NG: THE PIN ADDRESS OF THE REMOTE LCN MUST BE ENTERED BY SELECTING THE TARGET AT THE BOTTOM OF THE DISPLAY. AN ATTEMPT WILL BE MADE TO ACCESS THE VIEW USING THE CURRENT NODE NUMBER FOR THE LOCAL VIEW. IF THAT NODE IS NOT OPERATING, OR THE PIN ID IS IN ERROR IT WILL BE FLAGGED AS AN INVALID ENTITY NAME. SELECT AND ENTER A DIFFERENT NODE NUMBER, WHICH WILL BE USED FOR THE CURRENT VIEW, AND THEN REENTER THE REMOTE PIN ID.

TO CHANGE THE VIEWING NODE, SELECT AND ENTER A RUNNING NODE # USING THE TARGET AT THE BOTTOM RIGHT. ONCE ACCESS HAS BEEN MADE FOR REMOTE LCN ACCESS, ENTERING A NODE NUMBER THAT IS NOT IN THE RUN STATE WILL BE FLAGGED AS AN INVALID ENTITY NAME.

TO RETURN TO THE LOCAL LCN VIEW, SIMPLY HIT THE RETURN TO LOCAL VIEW TARGET. NOTE THAT ACCESS TO THE LOCAL LCN VIEW IS ALSO OBTAINABLE USING THE LOCAL LCN's PIN ADDRESS, IF IT HAS ONE.

CONFIRM LCN AND NODE USED FOR VIEW WITH PSDP ENTITY NAME ABOVE ENTER PIN TARGET!

34959

Figure 31 - LCNVIEWR help screen

LVRLOG Display

A version/revision logging function is available from the LVRLOG display that generates the following information:

- Each node on the LCN, and each card in each node excluding I/O.
- Each UCN connected to the LCN, each controller on the UCN, and each card in each controller.
- Each Data Hiway connected to the LCN, and each Box on the Data Hiway with its configured slots.

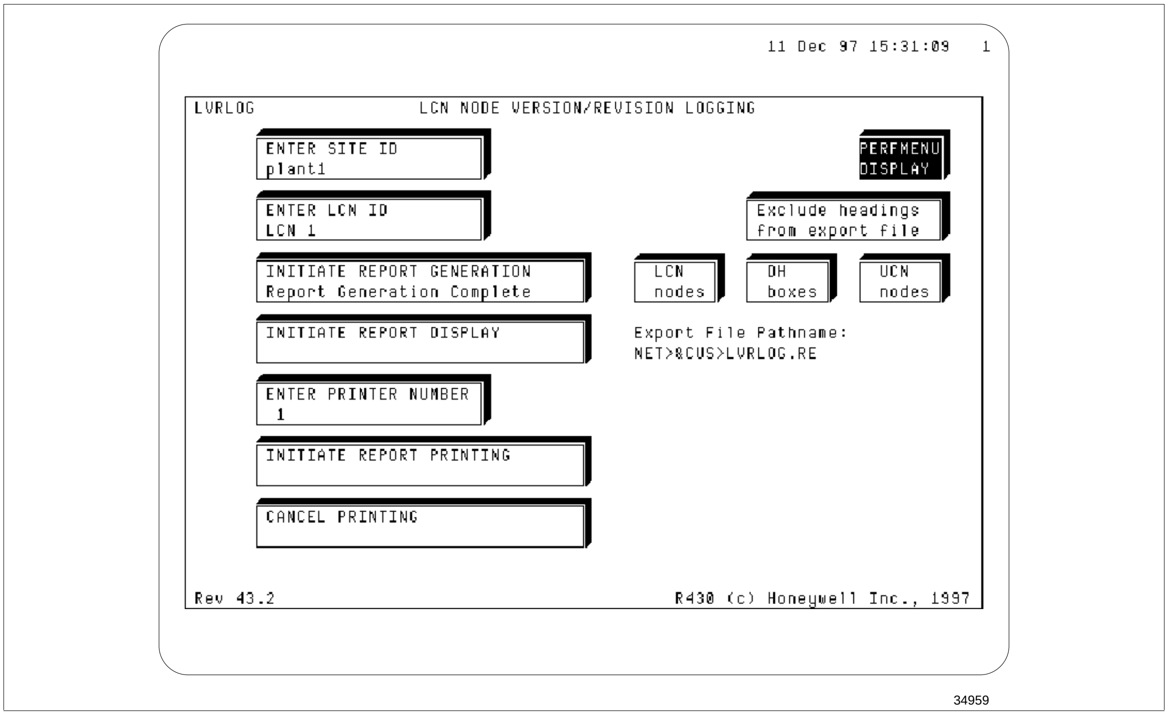

Figure 32 - LVRLOG Display

The information in the report includes the version and revision for the hardware, firmware, software currently being used on the system. This report includes information down to the card level. The limiting factor is the electronic accessibility of information from particular electronic modules (not all modules can be accessed for version/revision information). The report contains node, module, card type, and number information.

For the LCN nodes, the Version/Revision report includes

- the version and revision of the personality software currently loaded in the node,
- the name, version, revision, and the memory location of each External Load Module loaded in the node.

The information in the generated report is in text-file form and can be displayed or printed on demand. A text file with comma field separators is also generated for export to other applications such as Microsoft Excel. The export file path is:

NET>&CUS>LVRLOG.RE

The LCN Node Version/Revision Logging Function is an External Load Module (UPLVR or OPLVR) that resides in Universal Stations. This function is a TPS System backplane extension of the Universal Station Operator personality and the Universal personality.

HOLBRTH Display

This display indicates, for any particular node, its difficulty in accessing data from another node due to the other node’s loading.

11 Dec 97 15:32:11 1

HOLDBRTH - DISPLAY OF THE "HOLD BREATH" PARAMETER PFPSHLDB ON ALL NODES
HIGH VALUES OF PFPSHLDB INDICATE DIFFICULTY IN ACCESSING DATA IN
ANOTHER LCN NODE DUE TO THE OTHER NODE'S HIGH LOADING CONDITION R530

01	0	21	0	41	?????	61	?????
02	0	22	?????	42	?????	62	?????
03	0	23	0	43	?????	63	?????
04	0	24	?????	44	?????	64	?????
05	?????	25	?????	45	?????		
06	?????	26	?????	46	?????		
07	?????	27	?????	47	?????		
08	?????	28	?????	48	?????		
09	?????	29	?????	49	?????		
10	?????	30	?????	50	?????		
11	?????	31	?????	51	?????		
12	?????	32	?????	52	?????		
13	?????	33	?????	53	?????		
14	?????	34	?????	54	?????		
15	?????	35	?????	55	?????		
16	0	36	?????	56	?????		
17	?????	37	?????	57	?????		
18	?????	38	?????	58	?????		
19	0	39	?????	59	?????		
20	?????	40	?????	60	?????		

?????=NODE NOT LOADED

*****=VALUE > 99999

"PFPSHLDB" IS A DATA ACCESS PARAMETER ON EVERY NODE THAT INDICATES A HOLD BREATH EVENT COUNT SINCE LAST NODE LOAD OR VALUE RESET. THE COUNT IS INCREMENTED WHEN A TIME OUT IS VIOLATED FOR RECEIPT OF DATA FROM OTHER NODES ON THE LCN. THE TIME OUT VALUE DEPENDS ON THE FUNCTION LEVEL OF THE DATA ACCESS REQUEST AND RANGES FROM 1 SEC TO 20 SEC FOR MOST ON PROCESS REQUESTS. A FEW BUILDER TYPE REQUESTS AND NG REQUESTS RANGE FROM 25 SEC TO 90 MIN. USE DATACHNG TO RESET.

34956

Figure 33 - HOLBRTH Display

AM Load Balance Tool

This application analyzes AM points that are scheduled on the fast point processor only to provide a graphical representation of point loading per cycle on a specified AM. The graphical information is displayed in schematic $AMCYLD. This tool provides a composite set of results for all units and all cycles in the AM, and can also provide detailed information for all units on any given one-half second cycle.

In addition to the graphical representation of point loading, it also generates two text file reports:

1. AMxxLOAD.XX (where xx is the AM node number) which creates and analyzes an AM schedule dump file. The information within the file is displayed by cycle, unit by unit.
2. AMxxList.XX (where xx is the AM node number) which creates a list of points on a specific cycle across all units. It defaults to the most loaded cycle, but can be created for any cycle from schematic $AMCYLD.

The AM Load Balance Tool is not part of the standard TLK1 Displays, but is located in directory TLK2 for Release 600 and later. Each user must install this tool on their system; it does not come pre-installed.

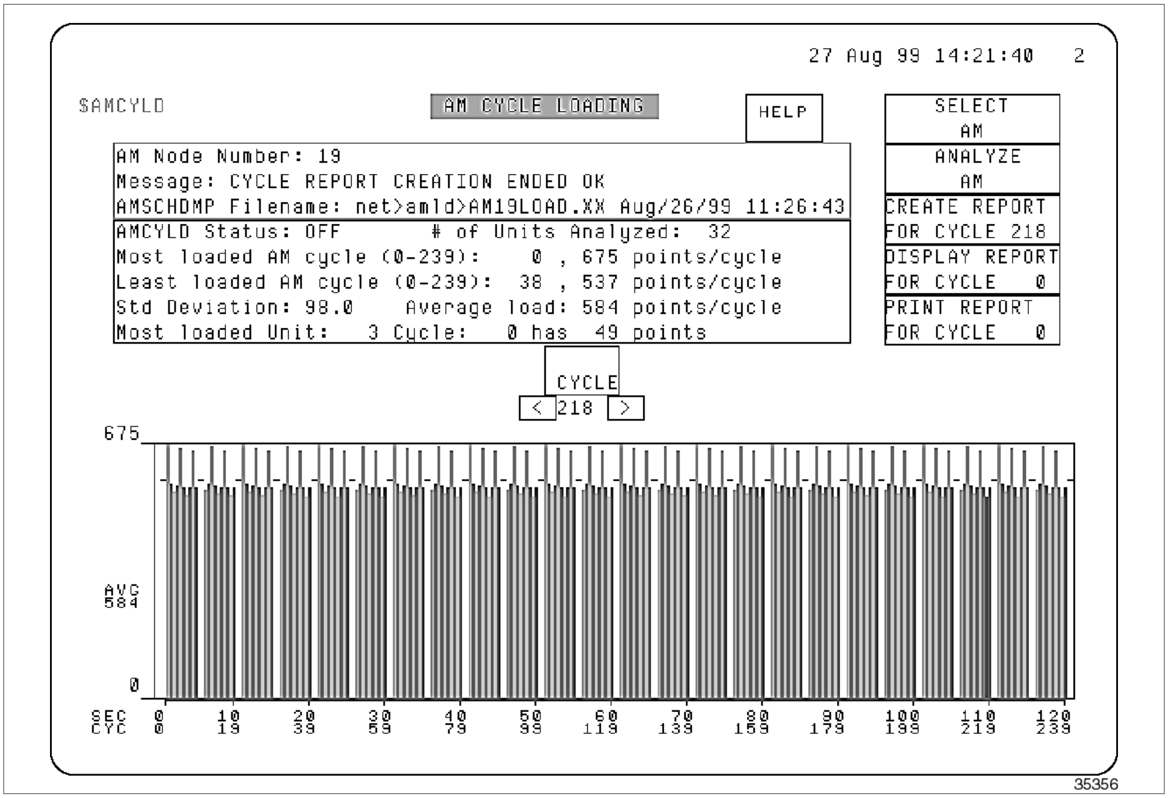

Figure 34 – AM Load Balance Tool

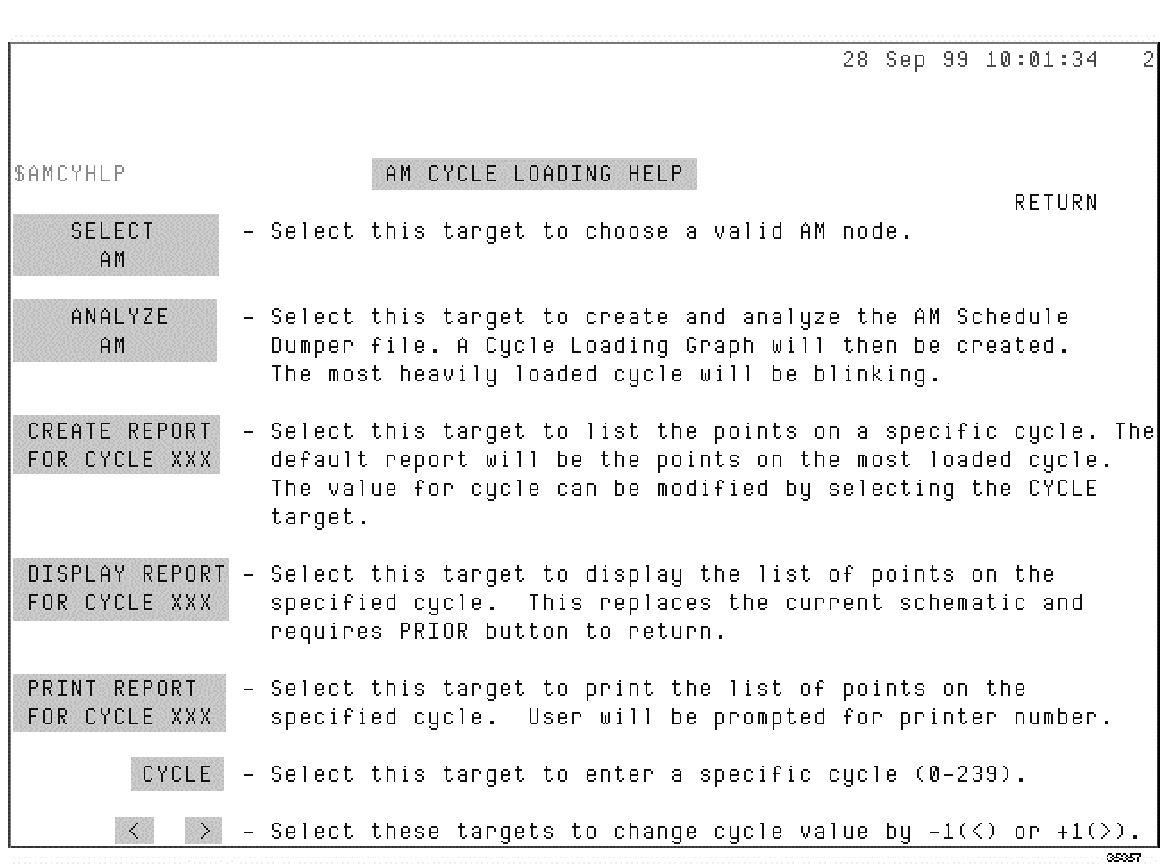

Figure 35 – AM Load Balance Tool Help Display

System-Specific Performance Displays

Description

The available standard performance displays are very useful; however, customized displays to support your specific system can provide a more convenient view into the system's operation.

Customized performance displays can be grouped by cluster, console, or Area, allowing convenient tracking if performance problems are detected or suspected across nodes.

PSDP Custom Trend Displays

Special trends of up to 16 parameters can be built to show performance characteristics of one or more nodes of interest on a console or cluster basis.

You may choose to trend the CPUFREE parameter for all nodes within the cluster. Another display can be built to trend the key PSDP parameters for the AM associated with that cluster.

These two displays would allow convenient tracking of the "culprits " when an event or action is reflected across several nodes, or loading is suspected within a node.

Suggested custom trends for a system are

- CPUFREE, EVTRATE, and PARSEC for all HGs/NIMs,
- CPUFREE, PARSEC, and GETAVGS for all AMs,
- GETAVGS, STRAVGS, CDSAVGS, and CLAVGS for all AMs, and
- NIM and HG CPUFREE and PARSEC with AM PFPSOVRC and PFPSHLDC.

Collecting History on PSDP Parameters

History cannot be directly collected on PSDP parameters. To historize data on these parameters you must use the AM.

> **ATTENTION**
>
> Caution must be taken when using the AM for gathering this kind of data because of the load that will be placed on the data owners. Response by any node for a request of PSDP parameters is handled as a high priority.
>
> Carefully watch your implementation to ensure that PSDP data gathering is not adversely affecting system performance.

Let us look at the case where the AM is requesting PSDP parameters from an HG. The HG's priority structure is as follows:

1. Manual output control
2. Alarming
3. AM output control/CG Access
4. US display call-ups
5. HM history collection
6. US display updates

Because PSDP requests are handled as a high priority, they compete with other high priority functions in the HG (such as manual output control and alarming). If the HG is very busy, excessive PSDP requests could interfere with the HG's ability to perform other high priority control activities.

> **ATTENTION**
>
> Any use of data resulting from the application of the displays, or the techniques described, to "your" system is totally your responsibility. Honeywell makes no claim as to the accuracy or usability of these displays and techniques for any purpose.

Baselining a TPS NETWORK System

Definition and Purpose

Baselining is providing a characteristic description of the system's performance under normal operating conditions. Given the baseline information, you can determine the general operational characteristics and growth potential of your system. Should a performance problem occur, this baseline information is invaluable in any investigation of the problem.

The purpose of baselining is to give you the information necessary to do the following:

- Determine when the system is reaching a physical limit under the current configuration,
- Determine the effects of increasing the existing database.
- Determine the effects of a software release upgrade.
- Troubleshoot problems, should they arise (such as slowing of display callups and HM or AM overruns).
- Determine your system's cycle. (Every system has a cycle unique to that system. The cycle is caused by activities such as history collection, AM data collection, and CG data collection.)

When to Baseline Your System

A baseline of your system should be performed at the following times:

- Once or twice a year when the system is under normal operation, and
- Before and after a release upgrade.

How to Use Baseline Data

If problems with a node or function occur, examine the conditions in the node or set of nodes supporting the function, then compare it with the baseline data to look for degradation.

Baselining Procedure

Purpose

Given the above data, the general operational characteristics and growth potential can be estimated.

If an when problems with a node or function occur, examine the condition of the node or set of nodes supporting the function, then compare it with the baseline data to look for degradation.

Should a performance problem occur, the baseline data will be invaluable in any investigation of the problem or in discussions with Honeywell TAC.

You should also consider the TPS Network performance monitoring service.

Procedure

Table 10 describes the baselining procedure. The procedure assumes that one or more Universal Stations with printers are available for the baseline activities. Customized displays for your system configuration can be very useful for this procedure.

Table 10 - Baselining Procedure

Step	Action
1	Over the period of an hour, make two or three screen prints of the following displays: • CPUCHKR • HEAPCHKR • HEAPFRAG • PARCHKR • HEAPMIN • CLOKSTAT
2	Obtain at least three trends covering 30-minutes of data for the HG, NIM, AM, HM, and NG: • QUIKTRND, • AMTREND, and/or • custom trends <u>Suggested Custom Trends for a system:</u> • CPUFREE, EVTRATE, AND PARSEC for all HGs/NIMs • CPUFREE, GETAVGS, AND PARSEC for all AMs • GETAVGS/STRAVGS/CDSAVGS/CLAVGS for all AMs • NIM or HG CPUFREE and PARSEC with AM PFPSOVRC and PFPSHLDC
3	Similar to the above step, obtain trends of the following for each US: • CPUFREE • NMSCHPR • HEAPMIN • AREAALM

Table 10 Baselining Procedure, (continued)

4	Over a one-hour period, take one or two sets of NODEPERF screen prints for each node on the system. **CAUTION** Do not leave NODPERF or any of the Detail displays on the screen—Print the display immediately after the data appears, then clear it from the screen (for example, call up the System Menu).
5	Over a one hour period, take one or two sets of the following screen prints for each AM, HM, and NG on the system: • AMDETAIL • HMDETAIL • NGDETAIL
6	Make one or two sets of the following screen prints for each UCN node: • UCNCOMM • UCN EVENT • UCNSUMM • NODESTA1/2
7	From the $LNMENU display, make screen prints of the following displays: • Cable A and Cable B Errors, • Reset Times, • NET Media Deads, • Small Buffer Status, and • Large Buffer Status.
8	Make a screen print of the following system displays: • System Status • any Custom System Status display • LCN Overview
9	Using SMCC, print the Module Errors Summary for each node on the system.

TPS NETWORK LCN Capacity Utilization Estimator

Description

Customers are concerned about the potential LCN load and capacity constraints; however, all tests indicate that for most customers the LCN load is not a major concern. All performance-related issues revolve around internal node capacity and/or design constraints; therefore, to address this issue, Honeywell is able to estimate the load on the LCN by totaling the parameters per second of all nodes on the LCN (use the PARCHKR display).

The LCN capacity estimator (illustrated in Figure 30) uses the PARSEC data available from the PSDP points in all nodes ($PRSTSnn.PARSEC).

There is currently no way to see the actual LCN loading in the TPS Network system. Extensive studies have been made of LCN loading on the Large System Test Facility (LSTF) in the Honeywell Customer Satisfaction Organization in Phoenix, Arizona. The estimator was empirically derived based on tests run on the LSTF.

> **CAUTION**
>
> The LCN estimator is provided as a guide. It may not cover all possible configuration or loading scenarios.
>
> Some standard displays (such as Group, Detail, and Overview) use “collection set” parameters that represent a much larger number of real parameters than is indicated on the performance displays; consequently, the LCN estimator should be used with care.
>
> Any decision made based on the estimator is not the responsibility of Honeywell.

How to Use the Estimator

Figure 36 illustrates the LCN estimator:

- X-Axis—The total PARSEC value.
- Y-axis—Indicates the % capacity load on the LCN.
- The linear line—A fixed progression that is used to determine the LCN load per the PARSEC value.

Use these guidelines when using the estimator:

- Smaller systems will have less base load because fewer nodes exist.
- A certain amount of access occurs just to keep the system running, even when there are no active nodes.
- Total all the PARSEC values of every node on the system. (If the PARCHKR display is available, simply call up this display and make a hard copy of it for future reference.)

Example

Given a total parameters per second (PARSEC) value of 1500, the corresponding estimated LCN load is 33% on a large system.

If the system was smaller with less history, slower checkpointing, and no events, subtract about 10% as indicated in the base load section to give a small system LCN load of about 23% when PARSEC total 1500.

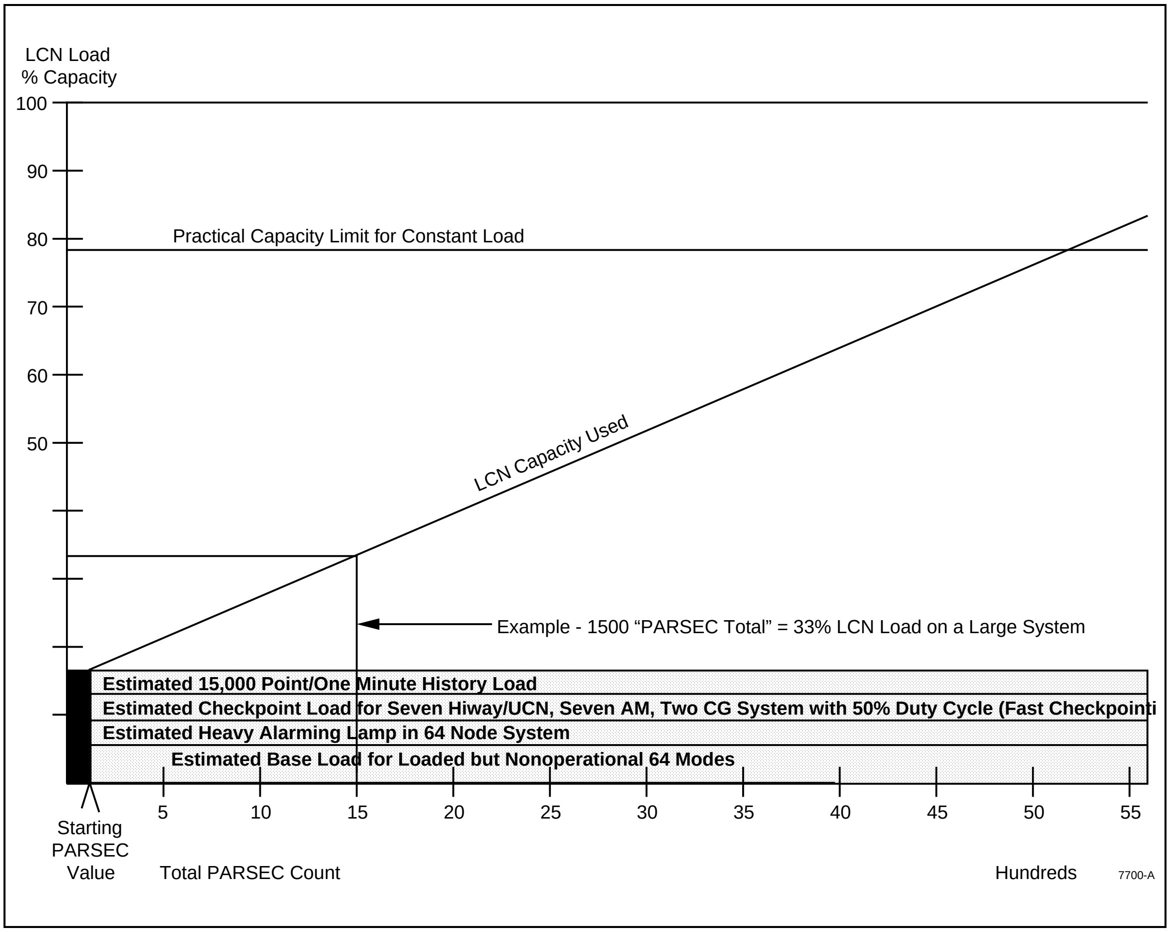

Figure 36 - LCN Capacity Utilization Estimator

Focused Load Concept

In relation to the performance load cluster, you must also consider the effect of a focused load on a node. A focused load occurs when a number of nodes request data simultaneously from a data owner and the amount of data requested is great.

Data is “owned” by such nodes as the HG, NIM, AM, and CG. When other nodes access a data owner to obtain data for such things as history and displays, a load is imposed on the data owning node.

Increasing the number of US/GUSs in a console, or temporarily increasing the number of US/GUSs using a display that accesses one or two data owners, during a plant upset, for example, can aggravate performance problems.

Remember that in the structure of the LCN, the typical hierarchy of access can cause a load to be focused on one data owning node, as the following figure shows, and can lead to performance problems on that node.

A focused load can also occur on a UCN node (such as a PM, APM, or HPM) when other nodes simultaneously request large amounts of data.

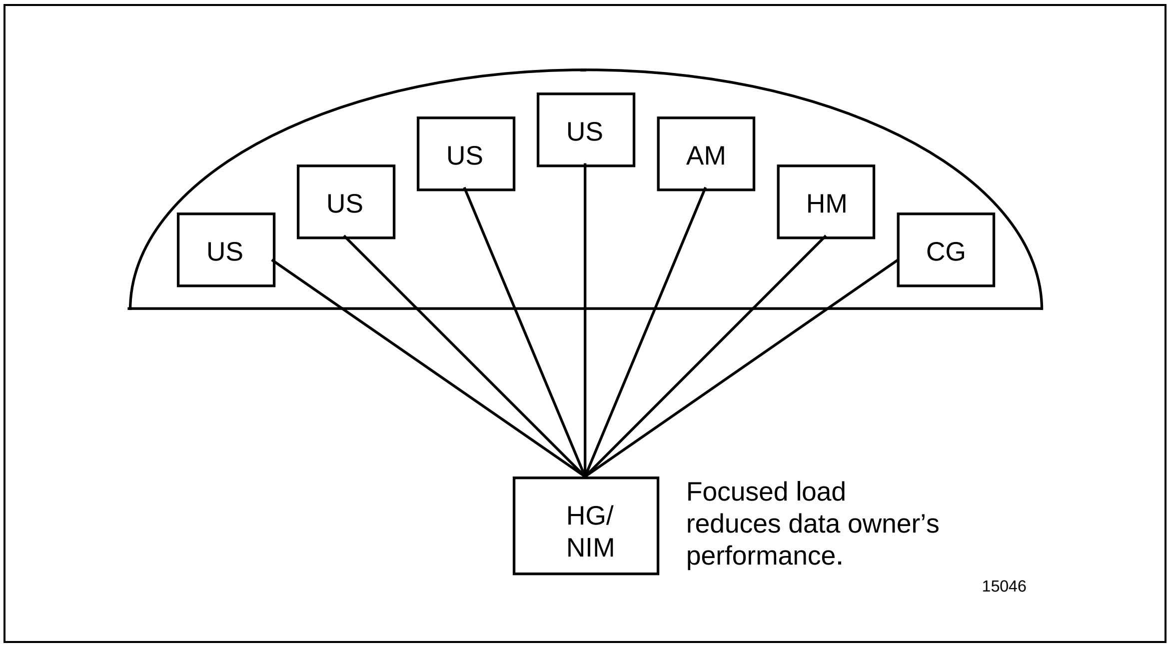

Figure 37 - Focused Load on Data Owner

LAB TIME

Use your assigned US/GUS.

Take with you:
- This module

LAB EXERCISE

The purpose of this lab exercise is to give you an understanding of how to use the displays listed in the PERFMENU display to determine the capacity of an LCN system.

1. Clear the screens all US/GUSs on the system.

   At US/GUS node n, build a 5-minute QUIKTRND display for US/GUS node m with the following parameters:

   RPARSEC

   RREQUEST

   PARSEC

   REQUESTS

   a. Call up a schematic at US/GUS node m.

      What parameters changed? Why?

2. Clear all US/GUSs on the system.

   At US/GUS node n build a 5-minute QUIKTRND display for US/GUS node m with the following parameters:

   CPUFREE

   RPARSEC

   PARSEC

   REQUESTS

   a. At US/GUS node m, call up the PIE2 display.

      (This display continuously redraws the display on the screen. The calculations behind it require all the CPU time.)

      What parameters changed? Why?

3. At US/GUS node n build a 5-minute QUIKTRND display for NIM node m with the following parameters:

   PARSEC

   REQUESTS

   RPARSEC

   SPARSEC

   a. Call up a schematic that uses points from that NIM.

      What parameters changed and why?

      Why don't we need to also plot RPARSEC AND SPARSEC?

4. At US/GUS node n build a 20-minute QUIKTRND display for HM node m with the following parameters:

   CPUFREE

   RPARSEC

   a. Can you tell from this display how often history is collected by this HM?

5. Using the PERFMENU display, familiarize yourself with the different performance displays, especially the checker, NODEPERF, and DETAIL displays.

   Capture a print of the LCN Overview display to determine the node types and their addresses on this system.

6. Using the PARSEC display and the LCN Estimator chart, determine the LCN capacity of this system.

End of Lab Exercise

Appendix A

Network Gateway Parameters

NG Details Display (NG Node Selected)

The following parameters are available from the NGDETAIL display.

TERM	Definition
CPUFREE	Percent of CPU free time
RPARSEC	Read parameters/second
PARSEC	Parameters/second
NGPAREQS	Parameter access requests
RECVWORD	Received words
RECVAVGM	Average rec.message size (bytes)
RECVMSGS	Messages received
RECVPKTS	512 word packets received
XMITWORD	Words transmitted
XMITAVGM	Average tx.message size (bytes)
XMITMSGS	Transmit messages
XMITPKTS	512 word packets transmitted
RECVFLMG	Receive file messages
RECVFLWD	Received message word count
XMITFLMG	Transmit file messages
XMITFLWD	Transmitted message word count
RECVERRS	Receive PIN errors
RECVTMOT	Receive time out
XMITERRS	Transmit PIN errors
XMITTMOT	Transmit time out

NG Details Display (Driver Detail Selected)

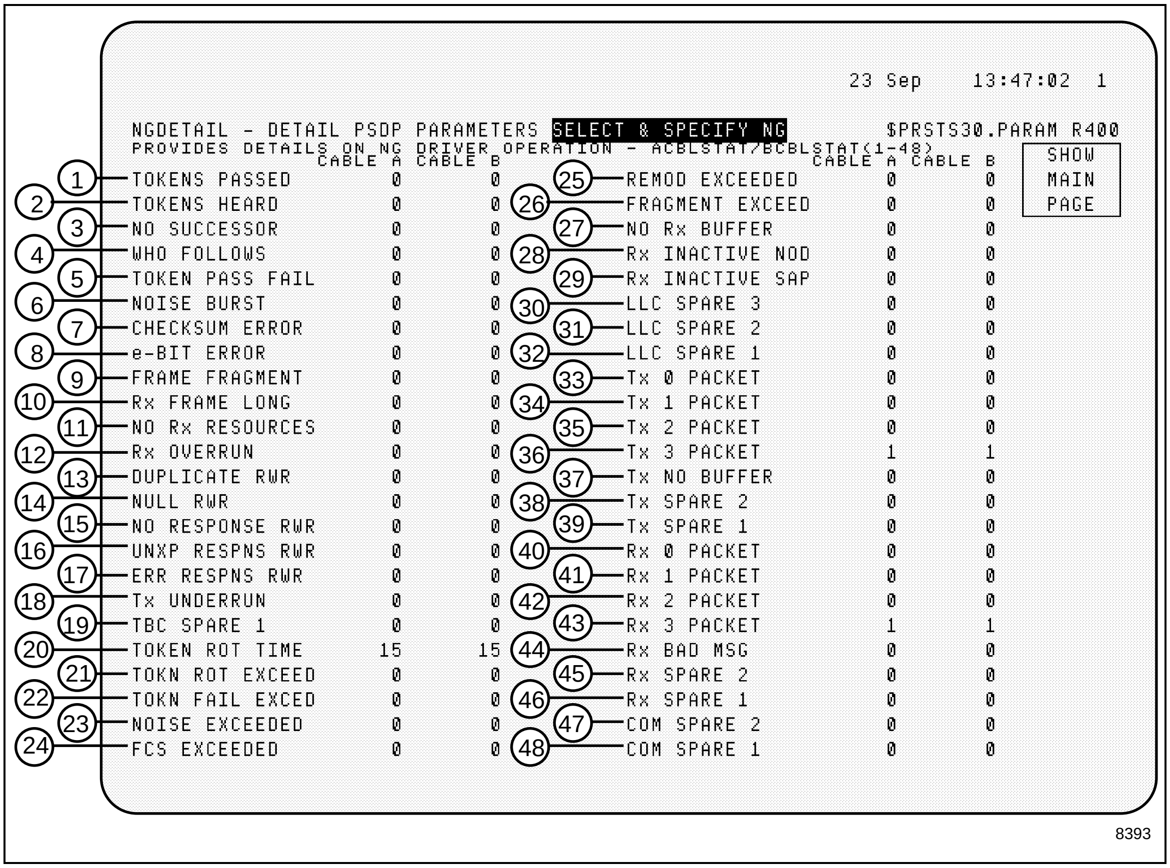

Network Gateway Statistics indicate values over the last minute. A list of acronyms used in the table is as follows:

Table 11 - Acronyms Used in Table 12

Acronym	Description
TBC	Token Bus Controller
DMA	Direct Memory Access
EMI	Electro Magnetic Interference
FIFO	First IN/First Out
LLC	Logic Link Control sublayer
RWR	Request With Response
SAP	Service Access Port
Rx	Receive
Tx	Transmit
FCS	Frame Check Sum

Table 12 - Term Definition for NG Driver Detail Display

Reference Number	Description
1	The TBC counts the number of tokens it successfully passes to its successor.
2	The TBC counts the number of tokens it hears that are not addressed to it. In a 2-node network this counter is zero. These first two statistics are zero unless enabled. Enabling them causes the TBC to perform a larger number of DMA cycles, which may have an impact on performance, Rx DMA overrun, and TxDMA overrun.
3	The TBC counts the number of times that, after being part of a ring, it could not find a successor. This means the ring collapsed. A very large value could indicate a faulty transmitter in this node. No_Successor normally is preceded by Who_Follows.
4	The TBC counts the number of times it had to look for a new node to pass the token to. Normally a count here is accompanied by a Token_Pass Failed count. If the TBC could not find a successor, Who_Follows is incremented by two and No_Successor is incremented by one.
5	The TBC counts the number of times it had to repeat a token pass. If the retry fails, the Who_Follows message is sent. A number of token pass failures during steady-state operation is an indication of a node or network problem. Token pass failures in one node accompanied by errors in its successor node indicate a probable receiver problem in the successor node. The token is passed in a ________pattern.
6	The TBC counts the number of noise bursts it sees. Noise bursts are periods of nonsilence that the TBC cannot interpret. The TBC cannot distinguish between noise caused by reflections and that caused by EMI. Noise_Bursts, Checksum_Errors, Remod_Errors, and Frame_Fragments are all considered noise errors.
7	The TBC counts the number of frame checksum errors that have the E-bit in the end delimiter reset.
8	The TBC counts the number of frames that have the E-bit in the End Delimiter set. The Head End Remodulator sets the E-bit whenever it receives a frame with a .checksum error. The error is propagated to the destination, with the E-bit set, for error tracking.
9	The TBC counts frames that start properly (preamble and Start Delimiter but do not end properly.
10	The TBC counts the number of received frames that are longer than the maximum of 8 KB.

(continued on next page)

Table 12 - Term Definition for NG Driver Detail Display (continued)

Reference Number	Description
11	The TBC counts the number of frames that were not received because there were not enough Frame Descriptors or Buffers.
12	The TBC counts the number of times the TBC detected a FIFO overrun (because of slow DMA) during receive. The complete message could not be copied.
13	The LLC counts the number of RWR frames with duplicate sequence numbers (Carrierband only).
14	The LLC counts the number of null messages. These are resynch messages to realign the message sequence numbers (Carrierband only).
15	The LLC counts the number of messages that did not get a response (Carrierband only).
16	The LLC counts the number of messages that contained the wrong MAC message type (Carrierband only)
17	The LLC counts the number of messages that contained a bad LLC message type (Carrierband only)
18	The LLC counts the number of times the TBC reported FIFO underrun (because of slow DMA) on transmit. The TBC could not complete transmission of the message.
19	The LLC counts the number of times the TBC reported that it was given a transmit request whose actual data length was greater than initially given to the TBC. The transmission is aborted.
20	The LLC samples the TBC token rotation time every 300 ms and performs a running average on it.
21	The LLC counts the number of times the TBC exceeded its preset Token_Rotation_Time threshold for a 300 ms period.
22.	The LLC counts the number of times the TBC exceeded its preset Token_Pass_Pass_Failed threshold for a 300 ms period.
23	The LLC counts the number of times the TBC exceeded its preset remodulation error threshold for a 300 ms period.

(continued on next page)

Table 12 - Term Definition for NG Driver Detail Display (continued)

Reference Number	Description
24	The LCC counts the number of times the TBC exceeded its preset FCS error threshold for a 300 ms period.
25	The LCC counts the number of times the TBC exceeded its preset remodulation error threshold for a 300 ms period.
26	The LCC counts the number of times the TBC exceeded its preset frame fragment error threshold for a 300 ms period.
27	The LLC counts the number of times it could not obtain a driver buffer to copy a message. Messages could be lost if driver buffers are not available for long enough for all LCC buffers to fill up.
28	The LLC counts the number of messages it discarded because the source node was not configured and SAP was not configured to accept all nodes. The overhead of copying an unwanted message is eliminated.
29	The LCC counts the number of messages it discarded because the destination SAP was not active. The overhead of copying an unwanted message is eliminated.
30-31-32	Spare (not used)
33-34-35-36	The driver counts the number of messages transmitted on each priority.
37	The driver counts the number of times a Get_TXB request could not be filled.
38-39	Spare (not used)
40-41-42-43	The driver counts the number of messages received on each priority.
44	The driver counts the number of unrecognizable received messages.
45-48	Spare (not used)

LAST PAGE