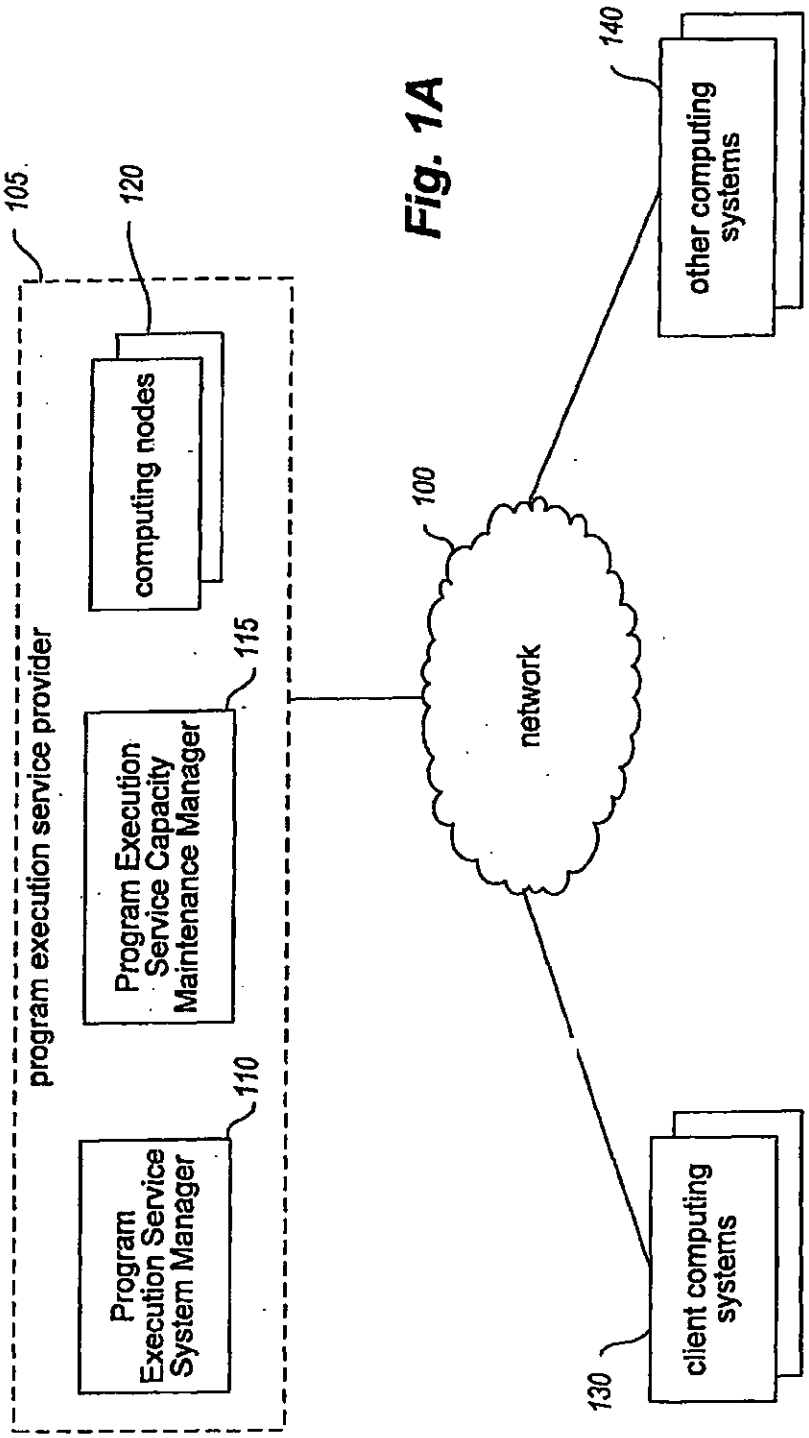
105
120
computing nodes
program execution service provider
115
Program Execution Service Capacity Maintenance Manager
110
Program Execution Service System Manager
100
network
140
other computing systems
130
client computing systems

Fig. 1A

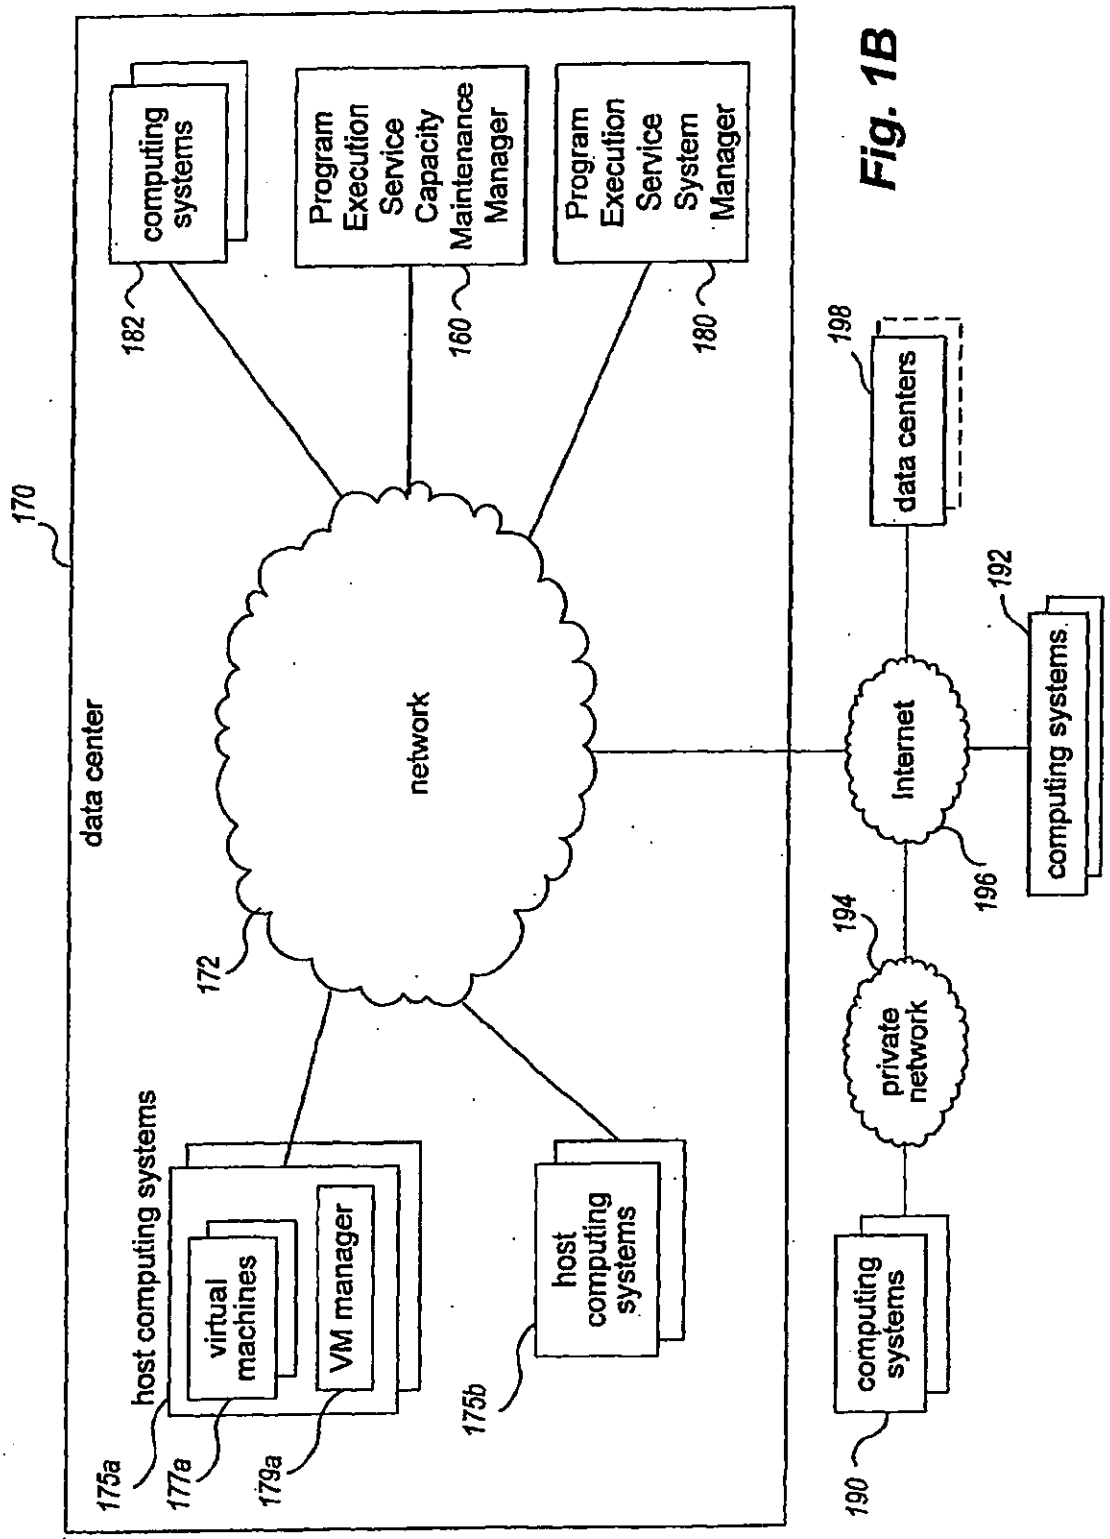
170
data center
computing systems
182
Program Execution Service Capacity Maintenance Manager
160
Program Execution Service System Manager
180
network
172
host computing systems
virtual machines
VM manager
175a
177a
179a
host computing systems
175b
data centers
198
Internet
196
computing systems
192
private network
194
computing systems
190

Fig. 1B

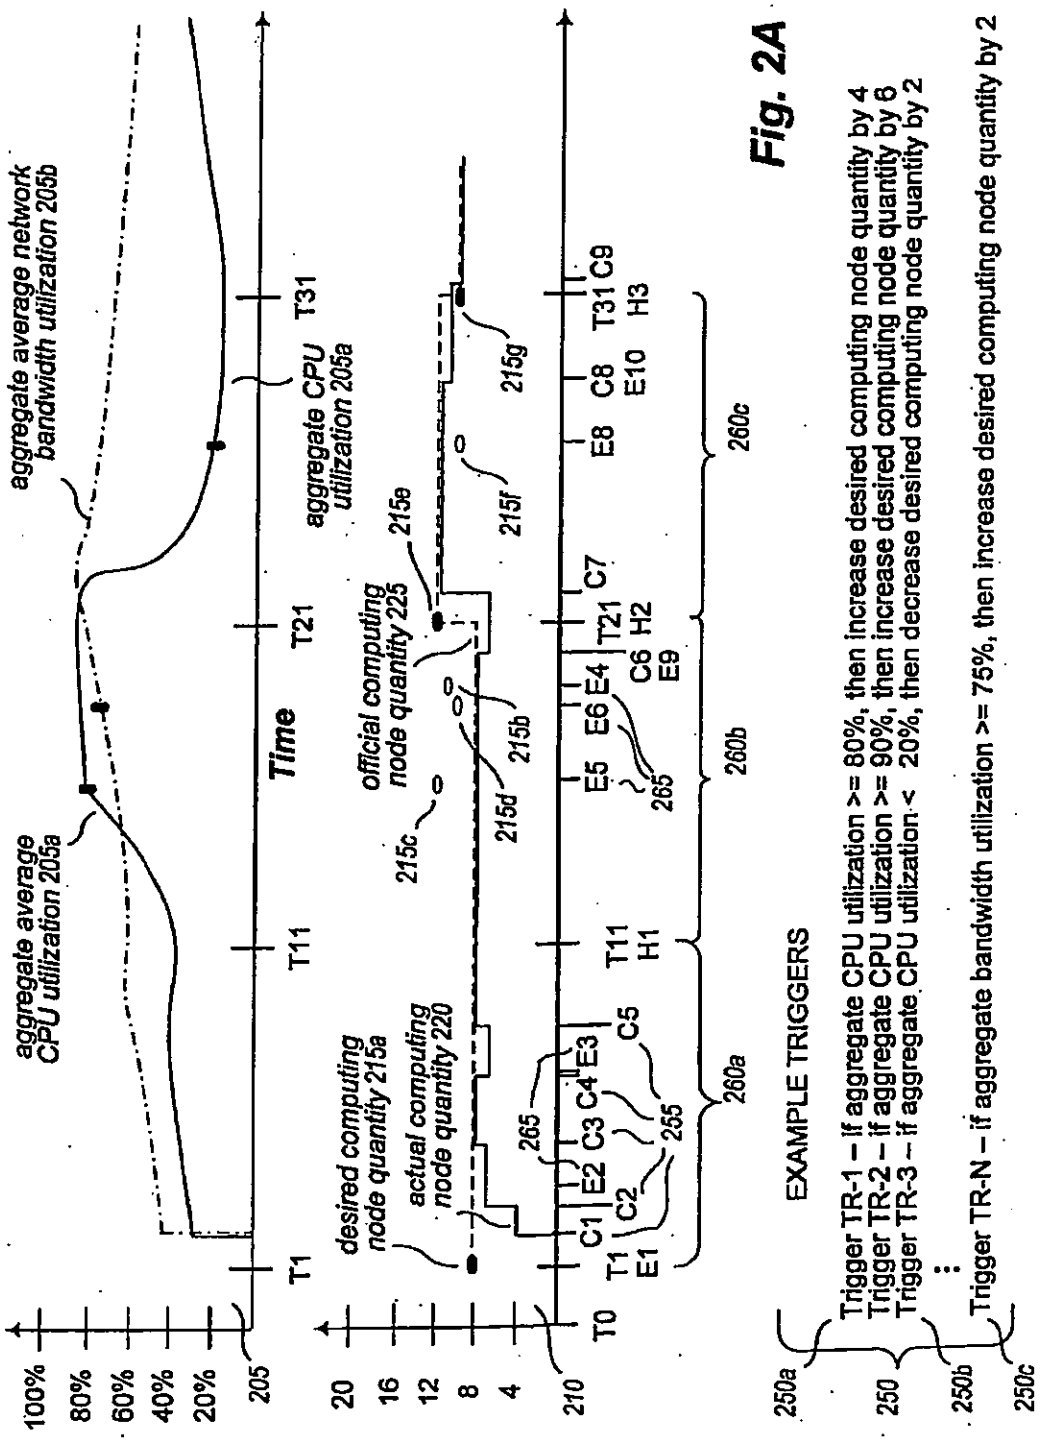
aggregate average network bandwidth utilization 205b
aggregate average CPU utilization 205a
100%
80%
60%
40%
20%
205
Time
T1
T11
T21
T31
aggregate CPU utilization 205a
official computing node quantity 225
desired computing node quantity 215a
actual computing node quantity 220
215c
215d
215b
215e
215f
215g
20
16
12
8
4
210
T0
T1
E1
C1
E2
C2
C3
C4
E3
265
255
C5
T11
H1
260a
E5
E6
E4
265
C6
E9
T21
C7
H2
260b
E8
C8
E10
T31
C9
H3
260c
EXAMPLE TRIGGERS
250a
Trigger TR-1 – if aggregate CPU utilization >= 80%, then increase desired computing node quantity by 4
Trigger TR-2 – if aggregate CPU utilization >= 90%, then increase desired computing node quantity by 6
Trigger TR-3 – if aggregate CPU utilization < 20%, then decrease desired computing node quantity by 2
250
250b
Trigger TR-N – if aggregate bandwidth utilization >= 75%, then increase desired computing node quantity by 2
250c

Fig. 2A

Example Computing Node Group Changes Database

280

	280a	280b	280c	280d	280e	280f	280g	280h		280x
	Change ID	Group ID	Change Type	Change Result	Change Details	Start Time	End Time	Time Period		Change-Event Nexus ID
285a	C01	Group1	computing node quantity	add nodes	+4 nodes	T01	T02	H1	...	N1
285b	C02	Group1	computing node quantity	add nodes	+3 nodes	T01	T03	H1		N1
285c	C10a	Group1	computing node quantity	add node	+1 node	T01	-	H1		N1
285d	C10b	Group1	computing node quantity	terminate uninitialized node	-1 node	T04	T04	H1		N2
285e	C03	Group1	computing node quantity	add node	+1 node	T04	T05	H1		N3
285f	C04	Group1	computing node quantity	terminate unavailable node	-1 node	T07	T07	H1		N4
285g	C05	Group1	computing node quantity	add node	+1 node	T07	T09	H1		N5
285h	C06	Group1	computing node quantity	terminate unavailable node	-1 node	T20	T20	H2		N6
285i	C07	Group1	computing node quantity	add nodes	+5 nodes	T21	T22	H3		N7
285j	C08	Group1	computing node quantity	terminate unavailable node	-1 node	T28	T28	H3		N8
285k	C09	Group1	computing node quantity	remove node	-1 node	T31	T32	H4		N9
	⋮									

Fig. 2B

Example Computing Node Group Events Database

290

	290a	290b	290c	290d	290e	290f		290x
	Event ID	Group ID	Event Type	Event Source	Time	Time Period		Change-Event Nexus ID
295a	E01	Group1	group initialization	user instruction	T01	H1	...	N1
295b	E02a	Group1	failed node initialization detected	system initiated	T04	H1		N2
295c	E02b	Group1	uninitialized node terminated	system initiated	T04	H1		N3
295d	E03a	Group1	unavailable node detected	system initiated	T07	H1		N4
295e	E03b	Group1	unavailable node terminated	system initiated	T07	H1		N5
295f	E05	Group1	increase desired quantity by 4	trigger TR-1	T16	H2		N7
295g	E06	Group1	increase desired quantity by 2	trigger TR-N	T18	H2		N7
295h	E04	Group1	increase desired quantity by 3	user instruction	T19	H2		N7
295i	E09a	Group1	unavailable node detected	system initiated	T20	H2		N6
295j	E09b	Group1	unavailable node terminated	system initiated	T20	H2		N7
295k	E08	Group1	decrease desired quantity by 2	trigger TR-3	T27	H3		N9
295l	E10a	Group1	unavailable node detected	system initiated	T28	H3		N8
295m	E10b	Group1	unavailable node terminated	system initiated	T28	H3		N9
	⋮							

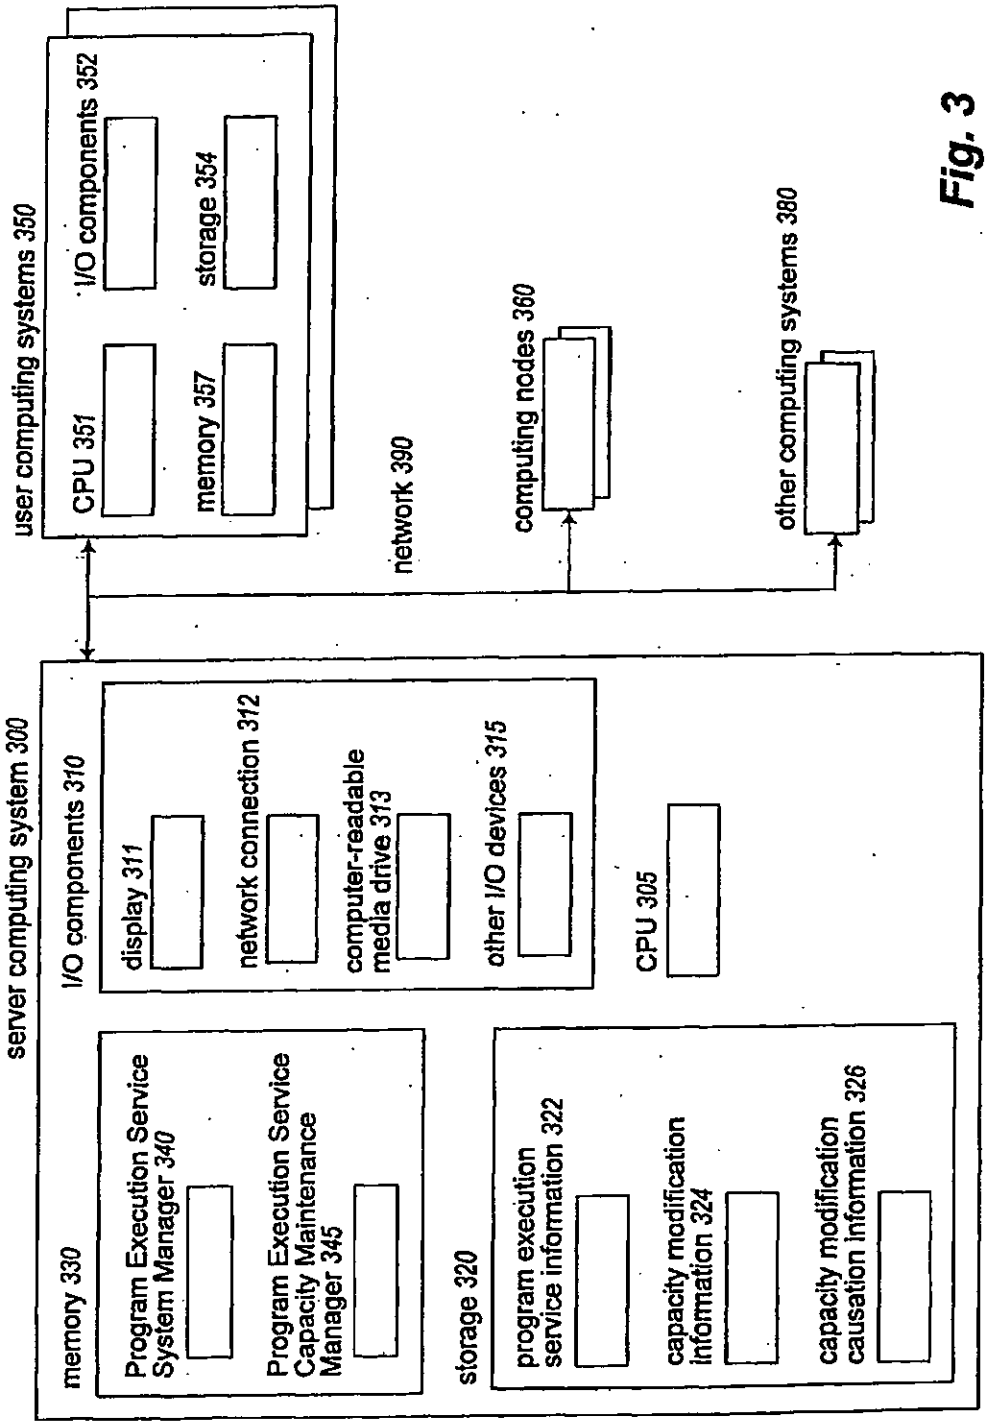
server computing system 300
memory 330
Program Execution Service System Manager 340
Program Execution Service Capacity Maintenance Manager 345
storage 320
program execution service information 322
capacity modification information 324
capacity modification causation information 326
I/O components 310
display 311
network connection 312
computer-readable media drive 313
other I/O devices 315
CPU 305
network 390
user computing systems 350
CPU 351
memory 357
I/O components 352
storage 354
computing nodes 360
other computing systems 380

Fig. 3

Fig. 4

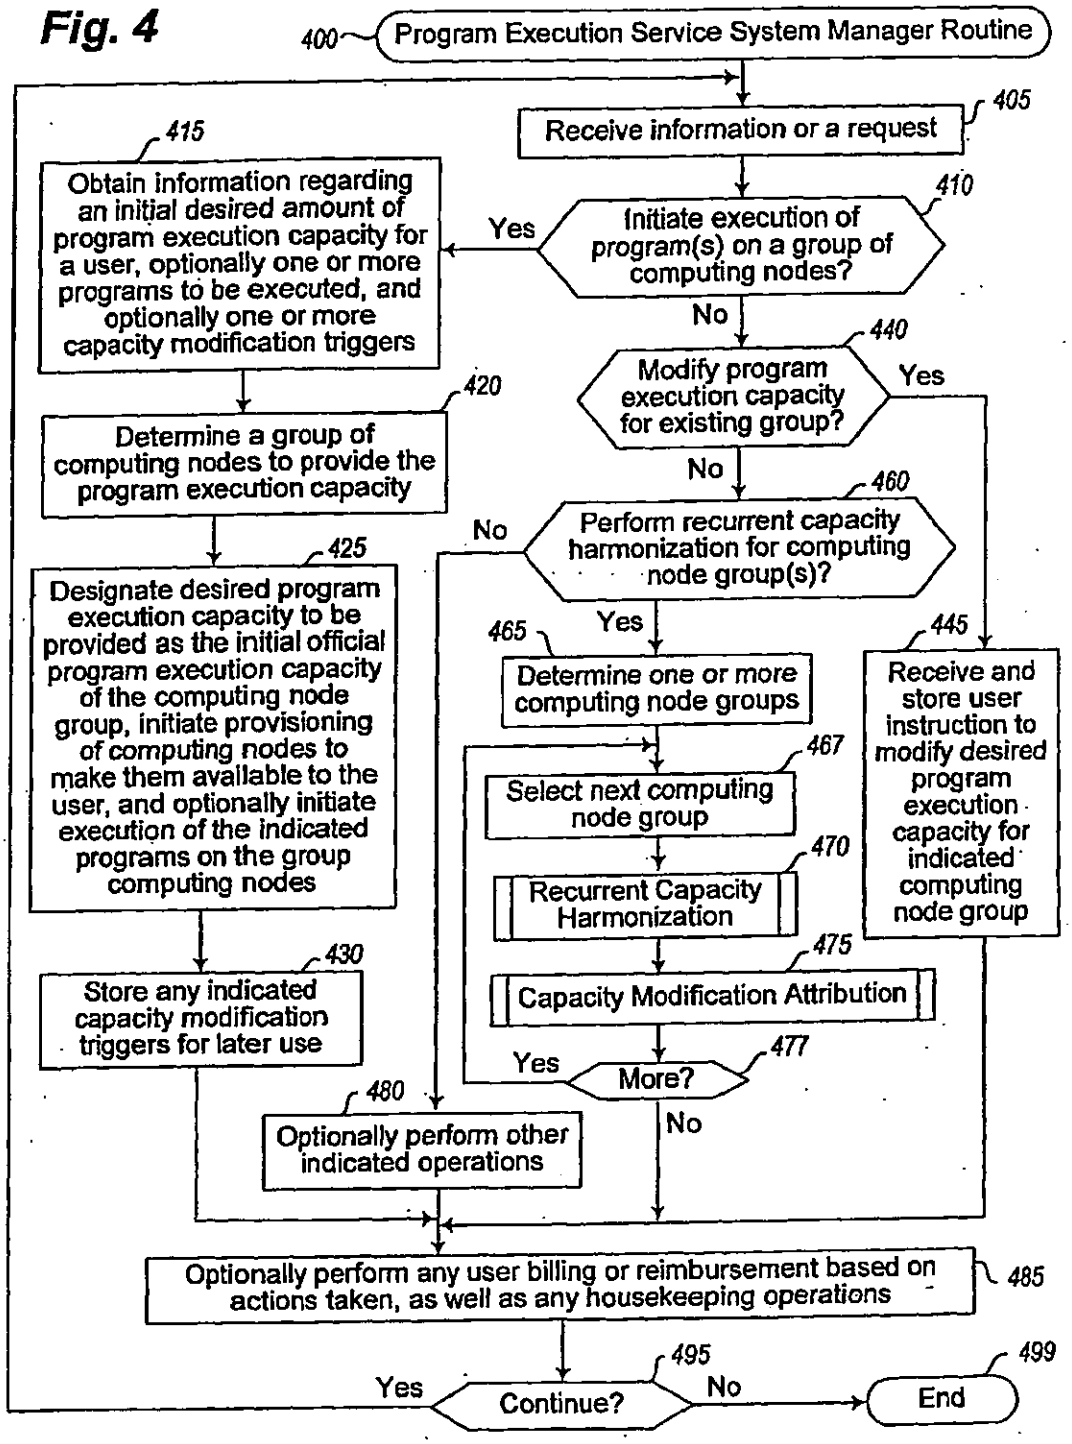
400
Program Execution Service System Manager Routine
405
Receive information or a request
410
Initiate execution of program(s) on a group of computing nodes?
Yes
415
Obtain information regarding an initial desired amount of program execution capacity for a user, optionally one or more programs to be executed, and optionally one or more capacity modification triggers
No
440
Modify program execution capacity for existing group?
Yes
420
Determine a group of computing nodes to provide the program execution capacity
No
460
Perform recurrent capacity harmonization for computing node group(s)?
No
425
Designate desired program execution capacity to be provided as the initial official program execution capacity of the computing node group, initiate provisioning of computing nodes to make them available to the user, and optionally initiate execution of the indicated programs on the group computing nodes
Yes
465
Determine one or more computing node groups
445
Receive and store user instruction to modify desired program execution capacity for indicated computing node group
467
Select next computing node group
470
Recurrent Capacity Harmonization
475
Capacity Modification Attribution
430
Store any indicated capacity modification triggers for later use
Yes
477
More?
No
480
Optionally perform other indicated operations
485
Optionally perform any user billing or reimbursement based on actions taken, as well as any housekeeping operations
495
Yes
Continue?
No
499
End

Fig. 5

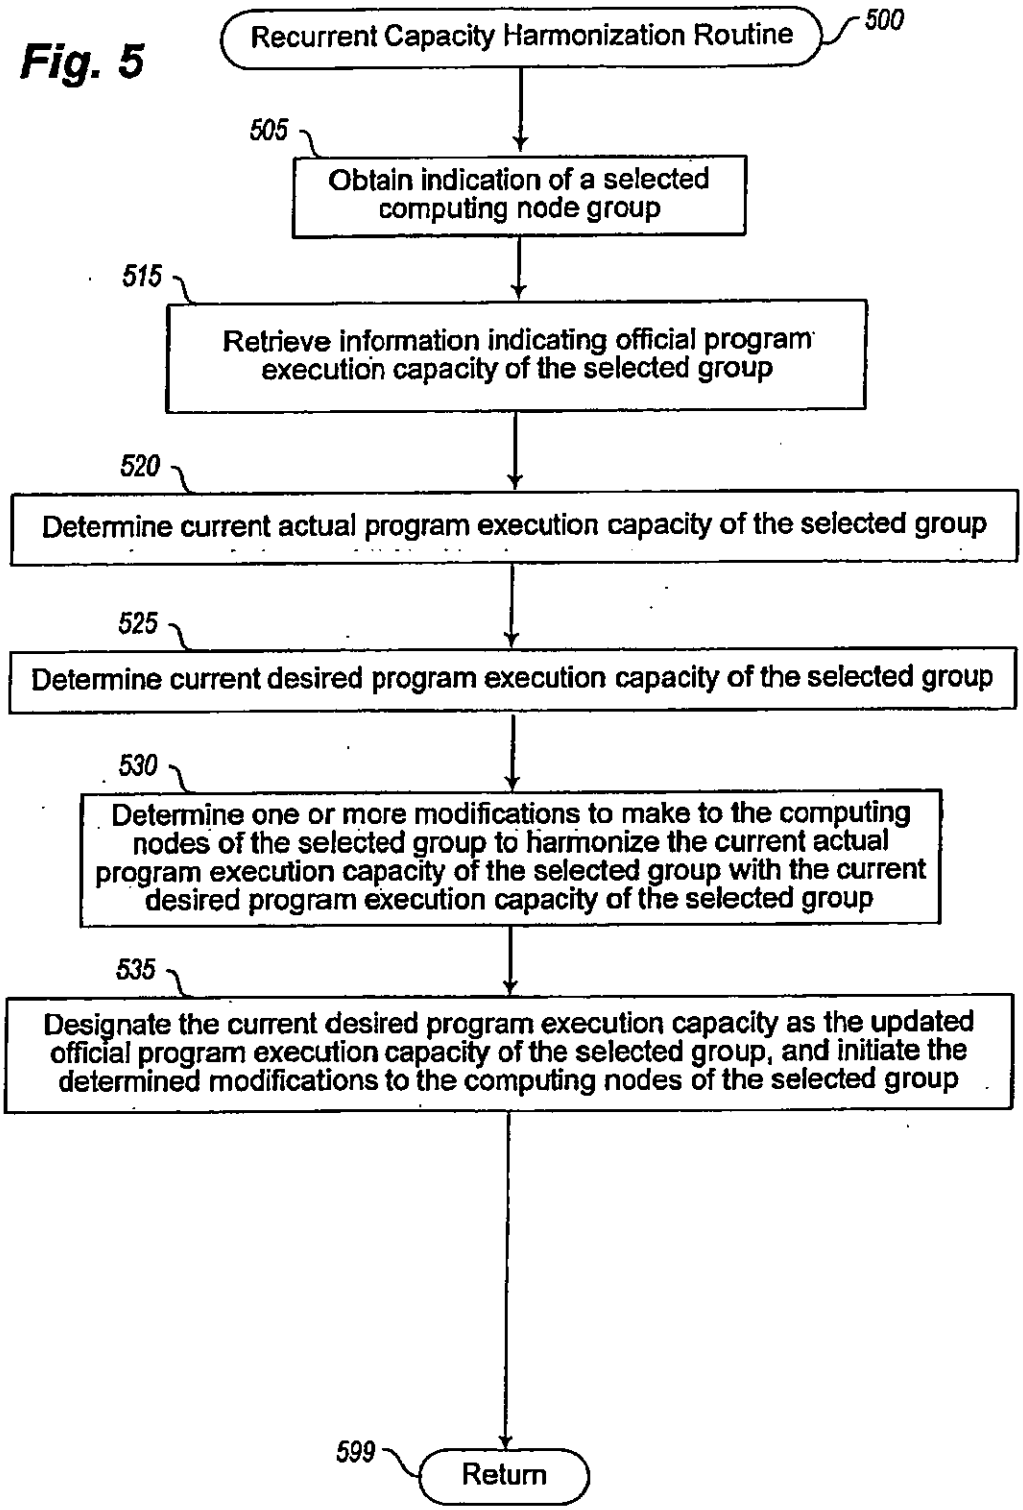
Recurrent Capacity Harmonization Routine
500
505
Obtain indication of a selected computing node group
515
Retrieve information indicating official program execution capacity of the selected group
520
Determine current actual program execution capacity of the selected group
525
Determine current desired program execution capacity of the selected group
530
Determine one or more modifications to make to the computing nodes of the selected group to harmonize the current actual program execution capacity of the selected group with the current desired program execution capacity of the selected group
535
Designate the current desired program execution capacity as the updated official program execution capacity of the selected group, and initiate the determined modifications to the computing nodes of the selected group
599
Return

Fig. 6

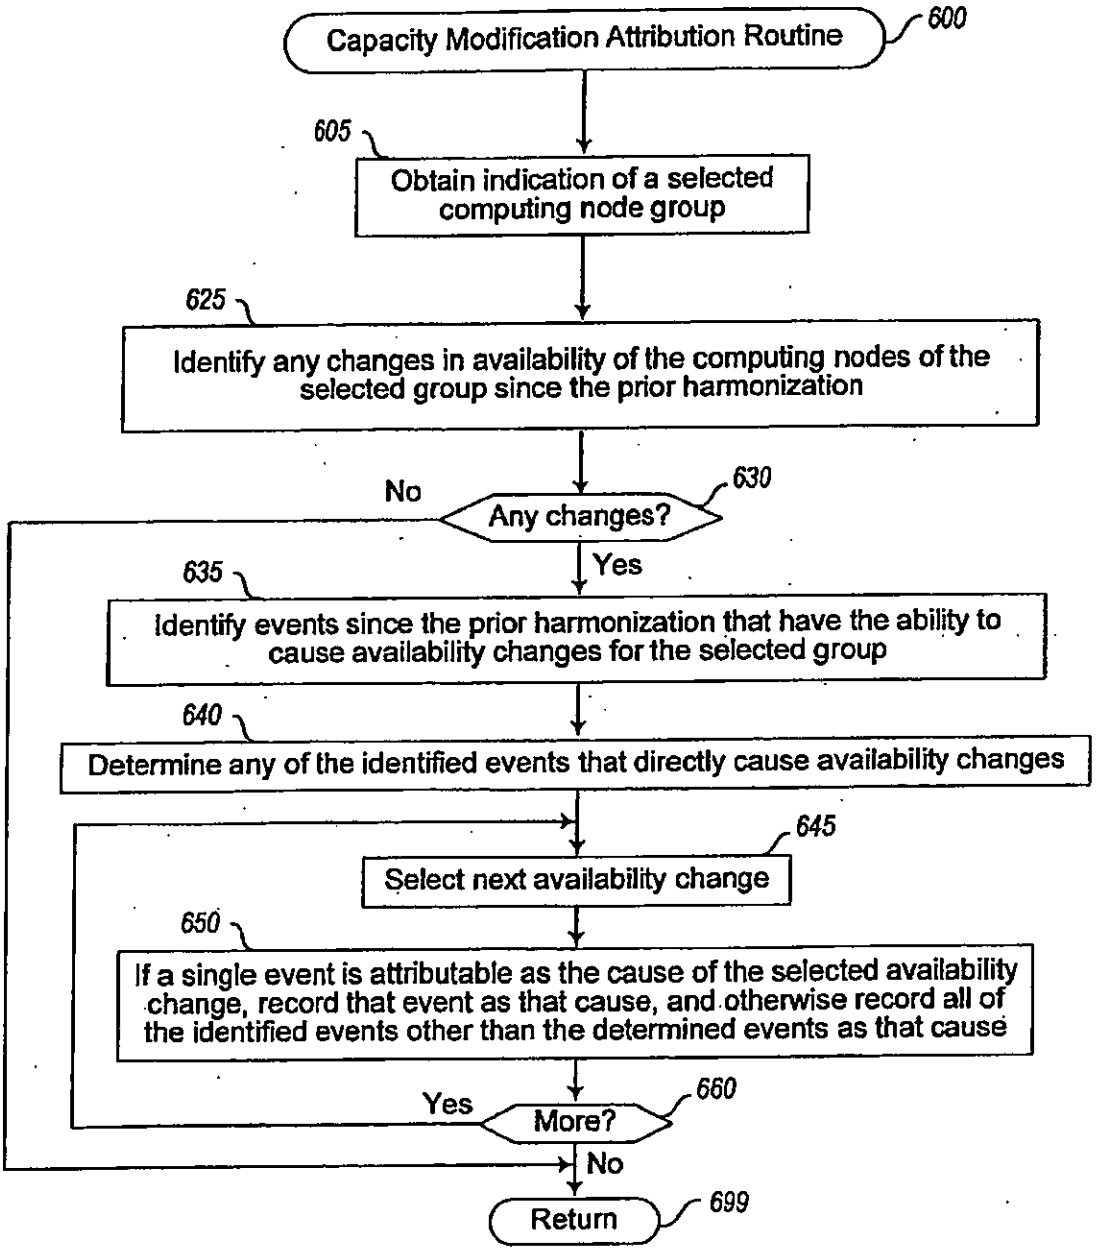
Capacity Modification Attribution Routine
600
605
Obtain indication of a selected computing node group
625
Identify any changes in availability of the computing nodes of the selected group since the prior harmonization
No
Any changes?
630
Yes
635
Identify events since the prior harmonization that have the ability to cause availability changes for the selected group
640
Determine any of the identified events that directly cause availability changes
645
Select next availability change
650
If a single event is attributable as the cause of the selected availability change, record that event as that cause, and otherwise record all of the identified events other than the determined events as that cause
Yes
More?
660
No
Return
699

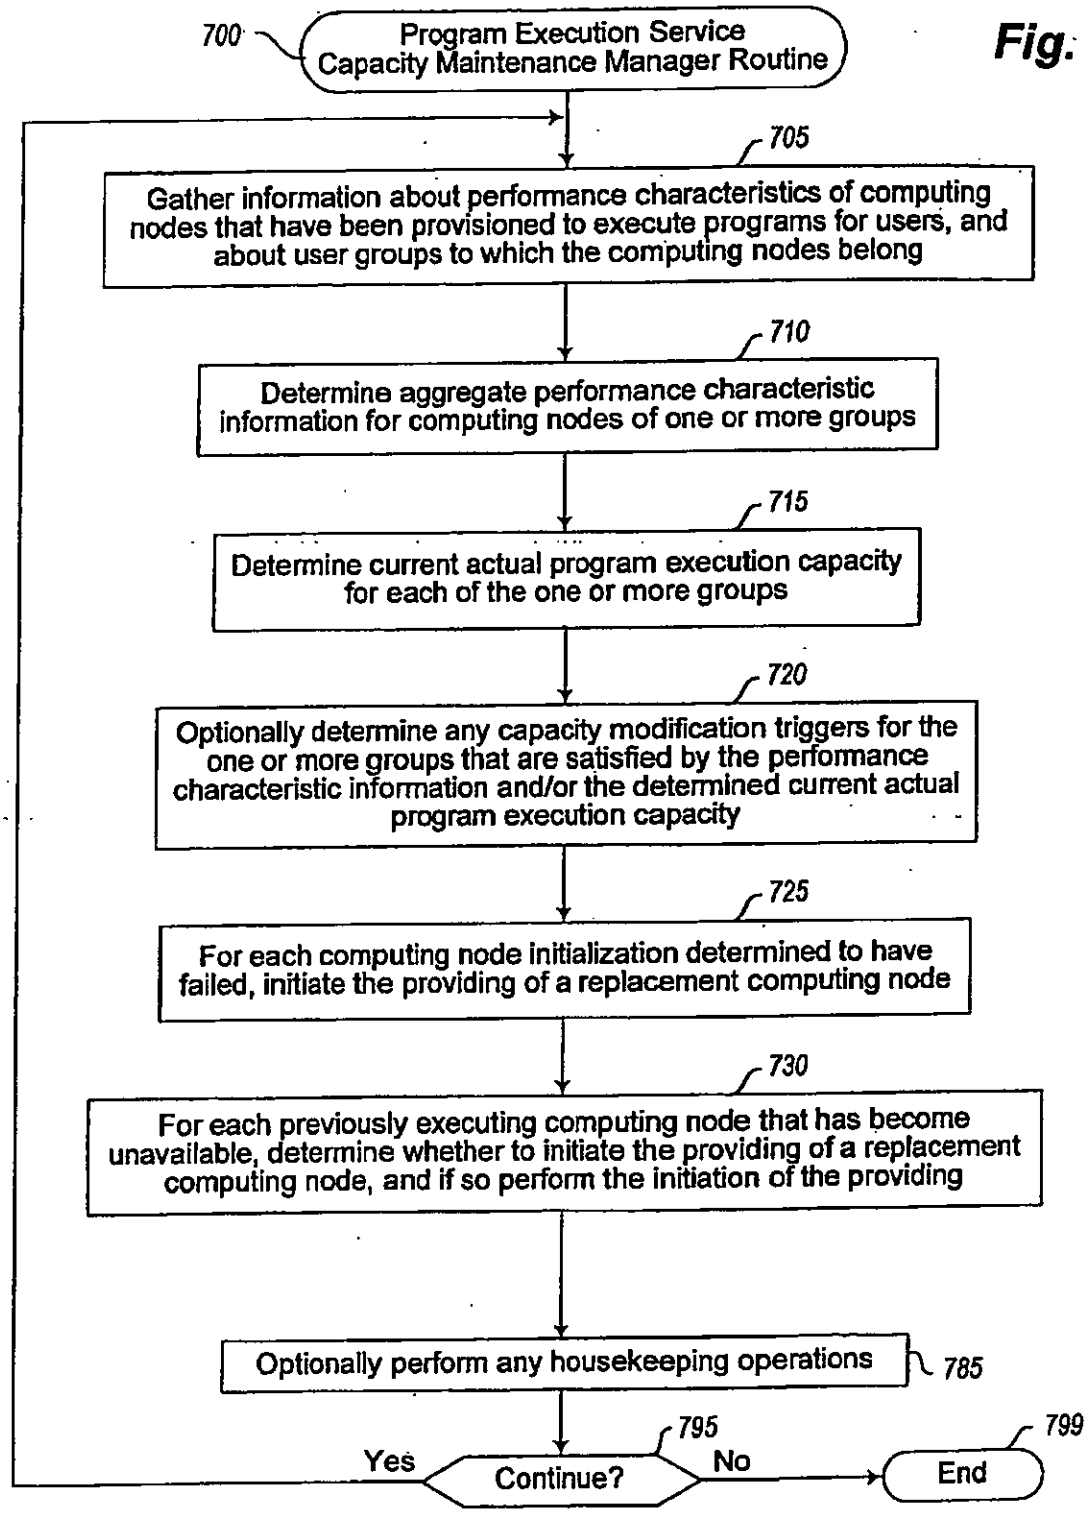
700
Program Execution Service Capacity Maintenance Manager Routine
705
Gather information about performance characteristics of computing nodes that have been provisioned to execute programs for users, and about user groups to which the computing nodes belong
710
Determine aggregate performance characteristic information for computing nodes of one or more groups
715
Determine current actual program execution capacity for each of the one or more groups
720
Optionally determine any capacity modification triggers for the one or more groups that are satisfied by the performance characteristic information and/or the determined current actual program execution capacity
725
For each computing node initialization determined to have failed, initiate the providing of a replacement computing node
730
For each previously executing computing node that has become unavailable, determine whether to initiate the providing of a replacement computing node, and if so perform the initiation of the providing
Optionally perform any housekeeping operations
785
795
Yes
Continue?
No
799
End

Fig. 7