

OpsDev – keeping performance and capacity in focus with development

Steven Thomas
SMT Data

November 2018
Session LM



DevOps and OpsDev

Both are a combination of development (Dev) and operations (Ops):

- **DevOps** focuses on making development more agile in deployment of new functionality to production
- **OpsDev** focuses on ensuring that operational considerations (performance, capacity, etc.) are in focus throughout development and deployment



1952 Model 41D
DeLuxe Touring Sedan





SMT Data invented ITBI™

TM Valuable new information for decision makers on all levels created through combined customer input regarding organisation, applications, finance, and selected technical IT-measurements -“BI with out-of-the-box value for IT”



We are devoted to enable large enterprises effectively run IT as a business, in order to save cost and align IT with business priorities



The largest companies are our customers

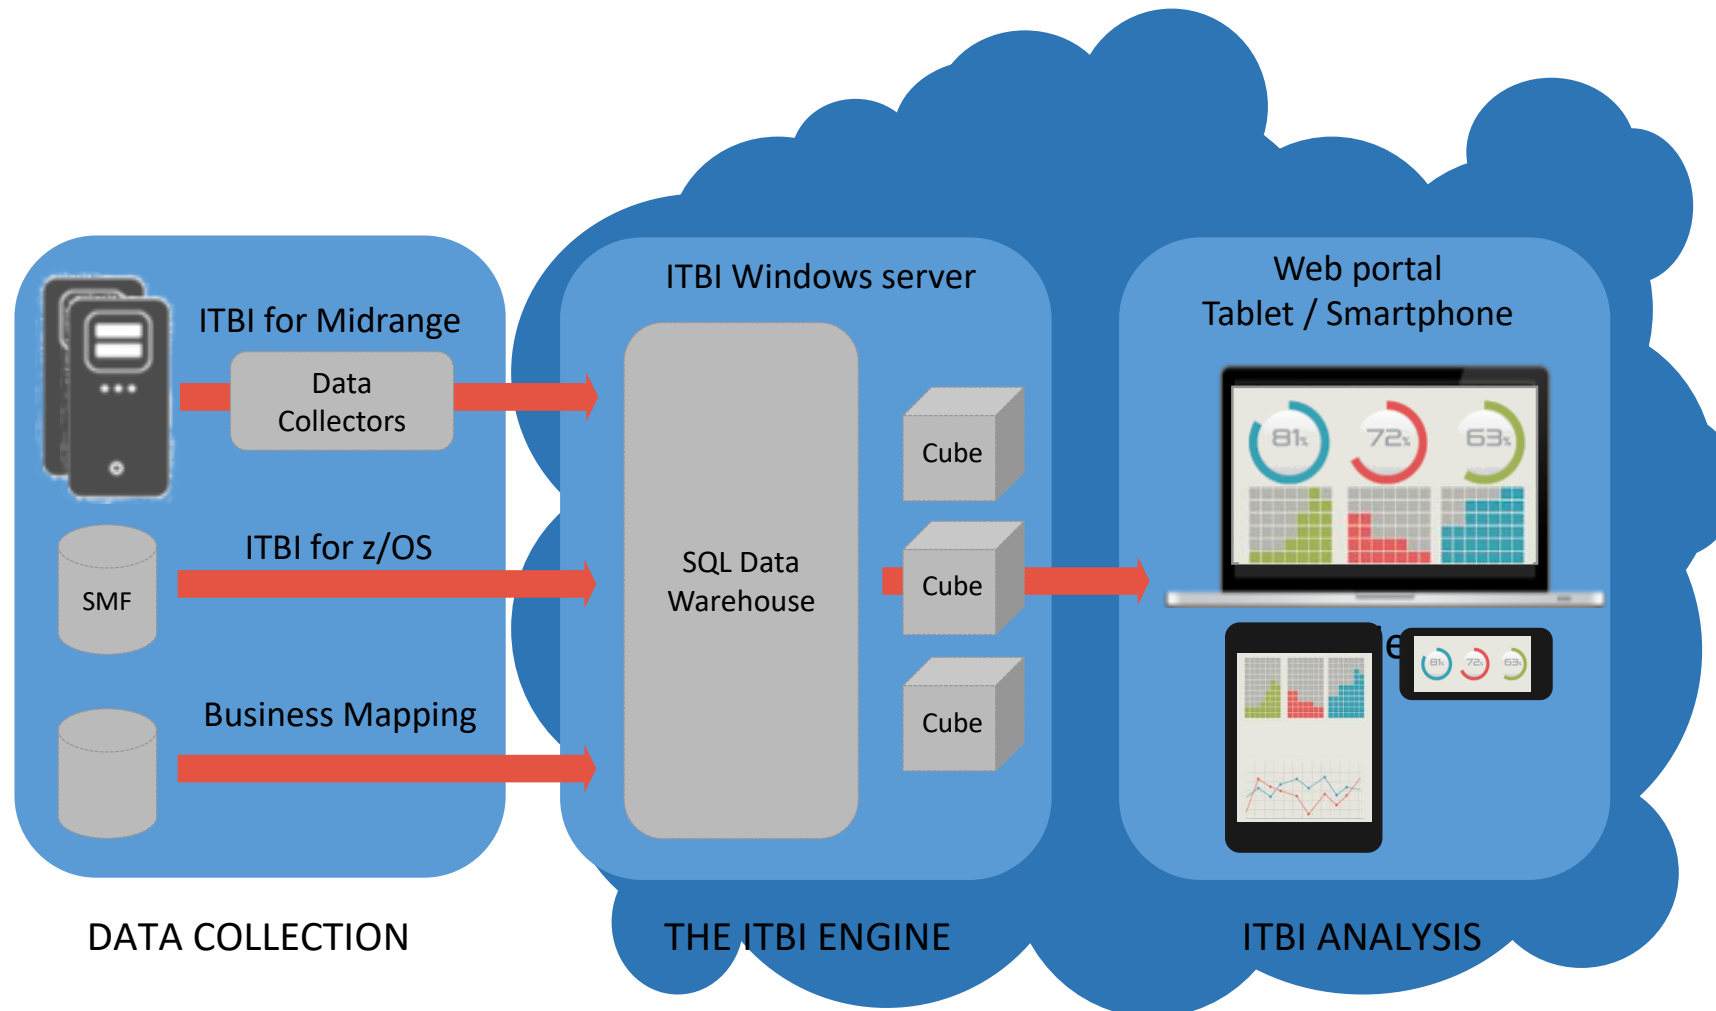


We team up with the leading international system integrators and other value adding partners to leverage the full value of ITBI



Our vision is to be the ITBI leader world wide


ITBI - Product architecture



OpsDev means measurement and feedback

Ongoing measurement and feedback from the IT-operations teams to the software developers:

- Identify candidates for tuning and prioritize these efforts alongside functional changes based on the financial impact
- Understand the impact of changes
- Tune and right-size the infrastructure
- Explain the capacity costs in business terms




Identify Candidates for Tuning


Identify Applications for Tuning - Mainframe

- Identify peaks
- Drill down to understand cost drivers e.g: LPAR, workload, job, transaction, DB2 package, user, etc.


Top-10 MIPS-Users for Selected Period




Hourly MIPS per Workload



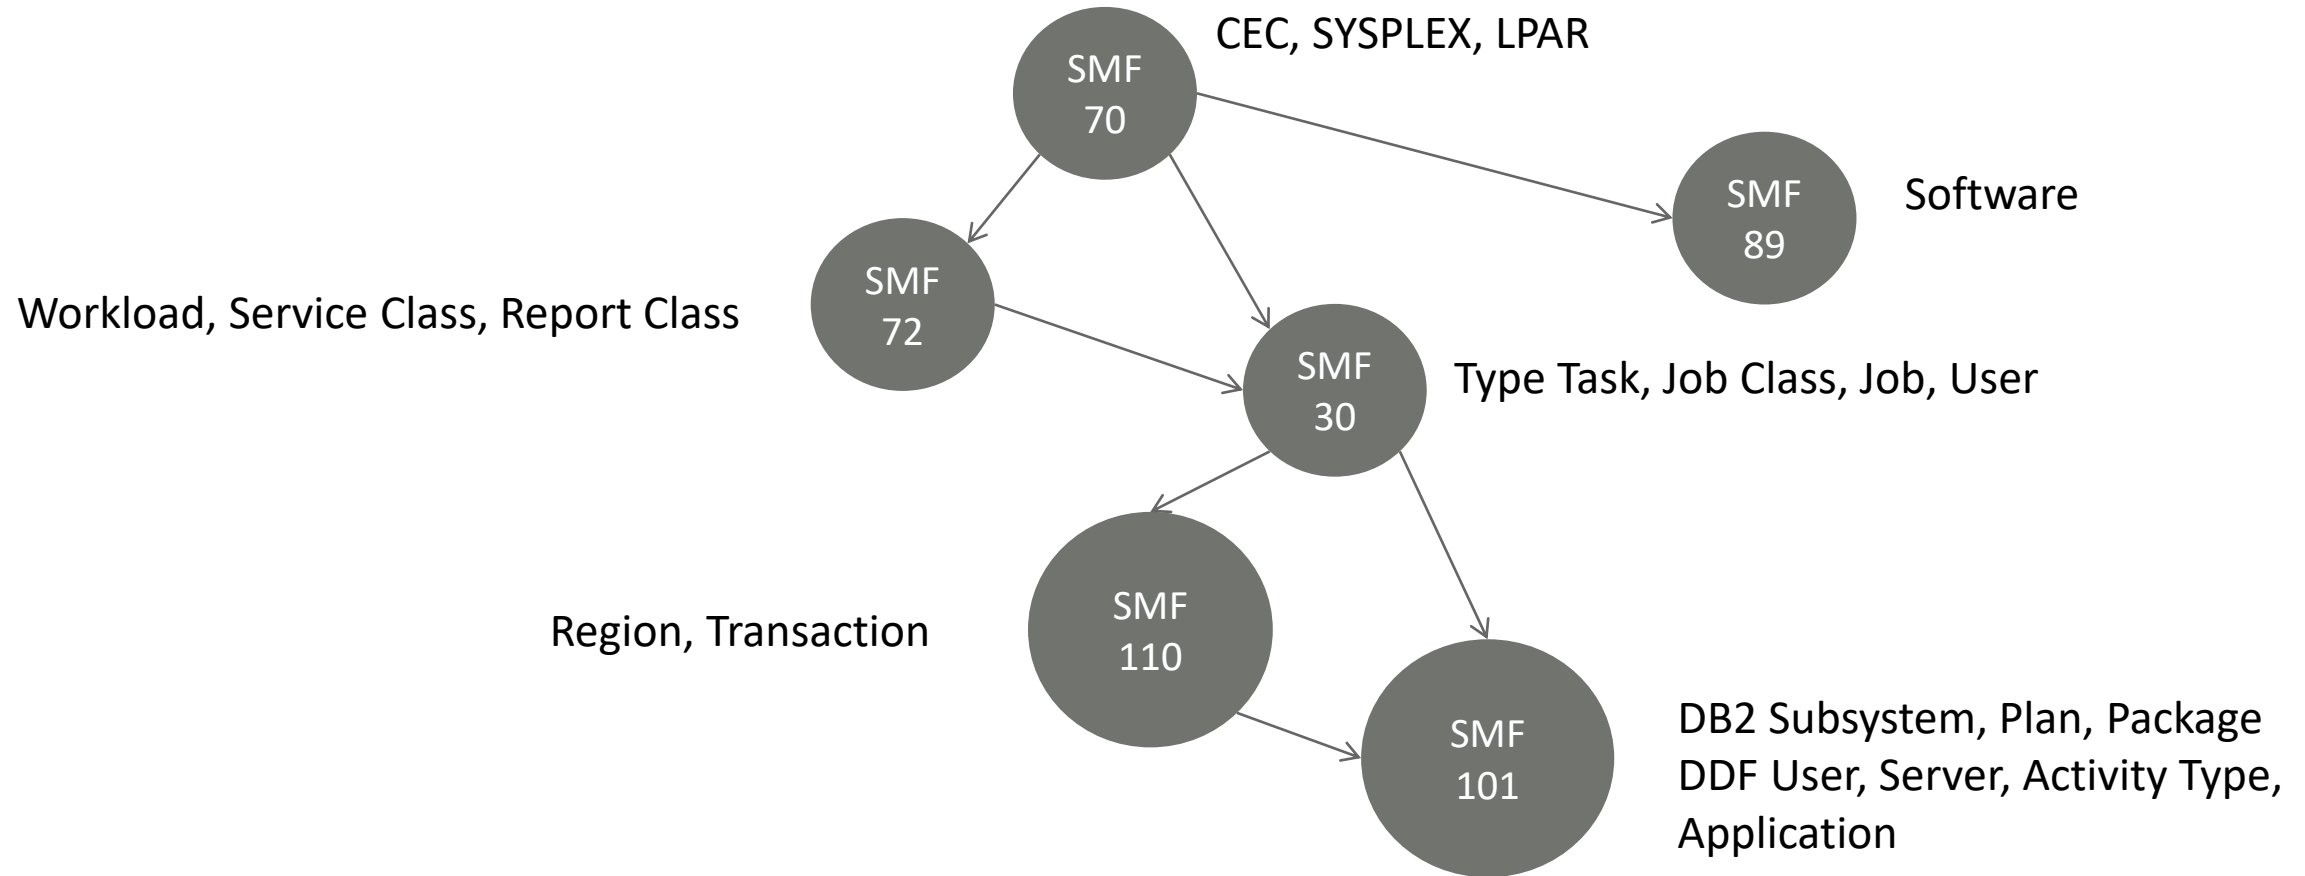
Hourly R4HA MSU Usage per Measured System



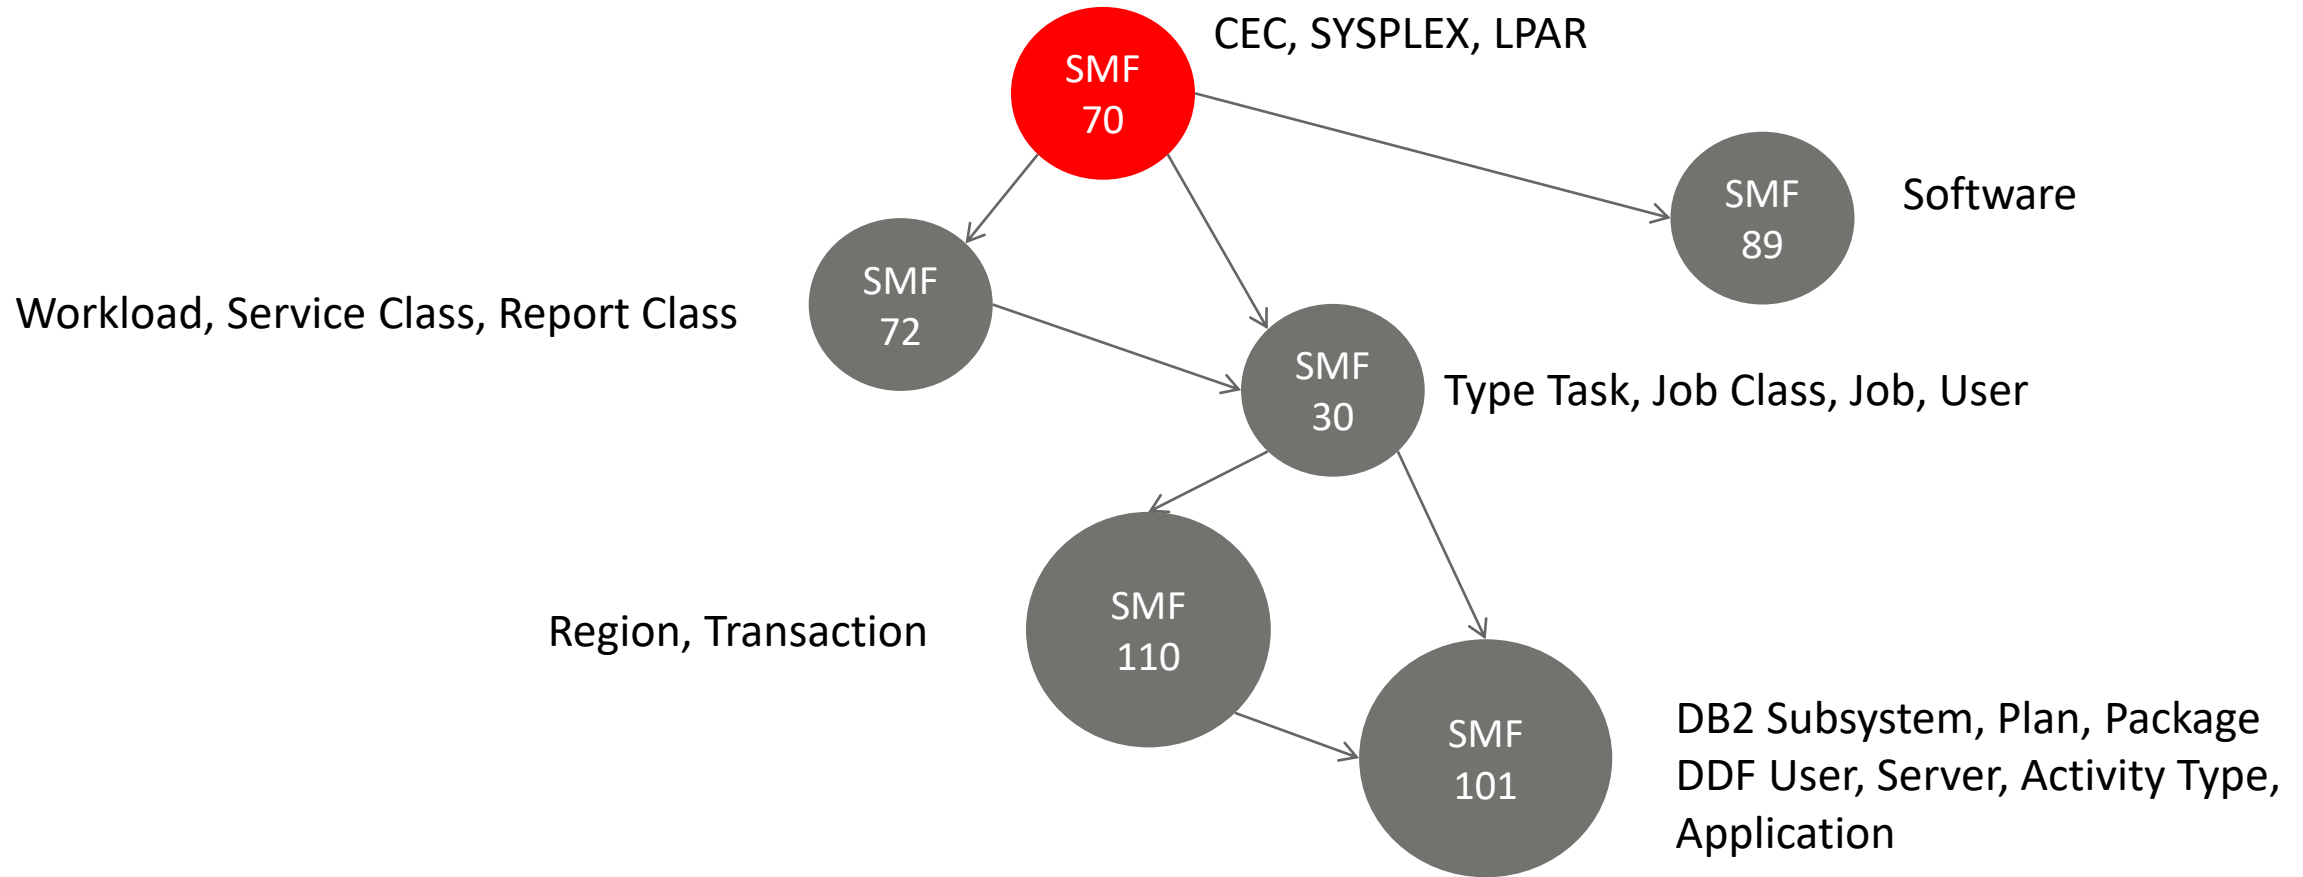
Hourly MIPS per Workload




Drilling down to understand peaks



Example: Test LPAR peak during month peak

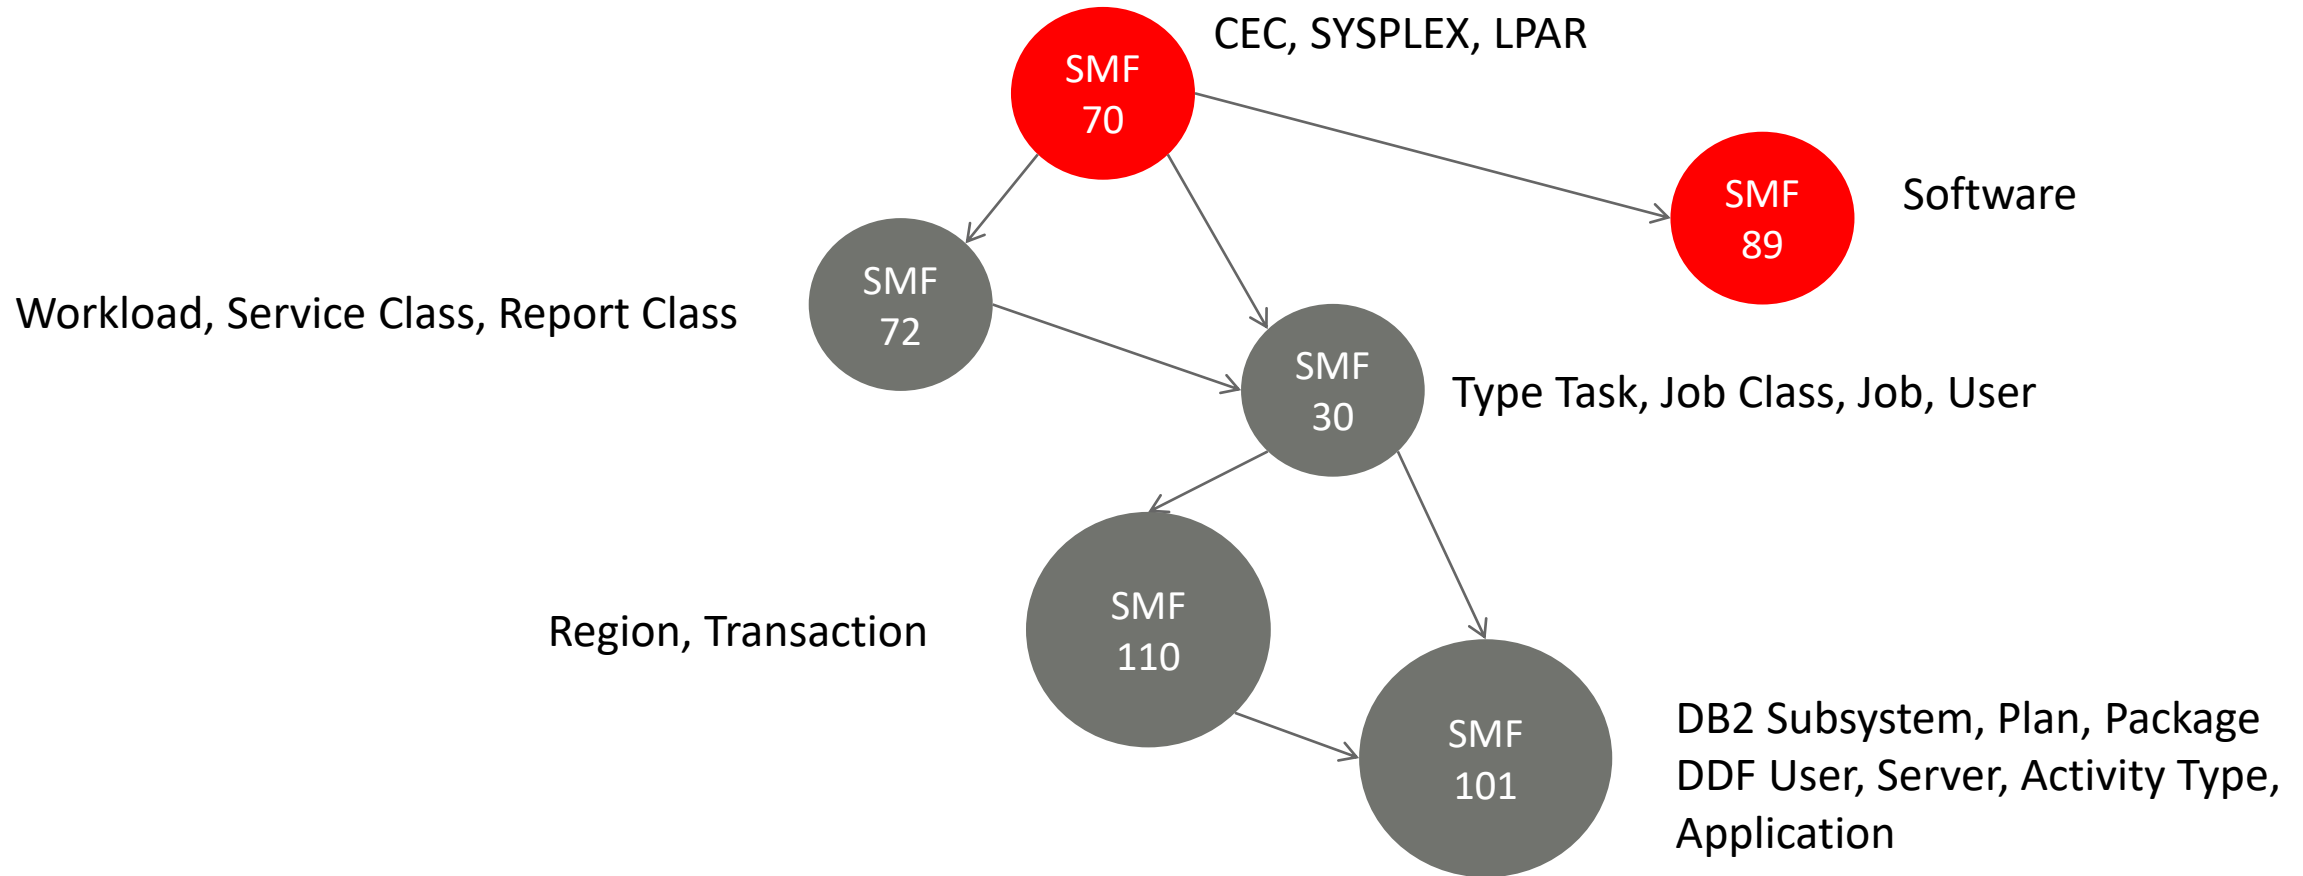


SMF70




Test LPAR using 300 MIPS during monthly peak online hours.

Example: Test LPAR using software



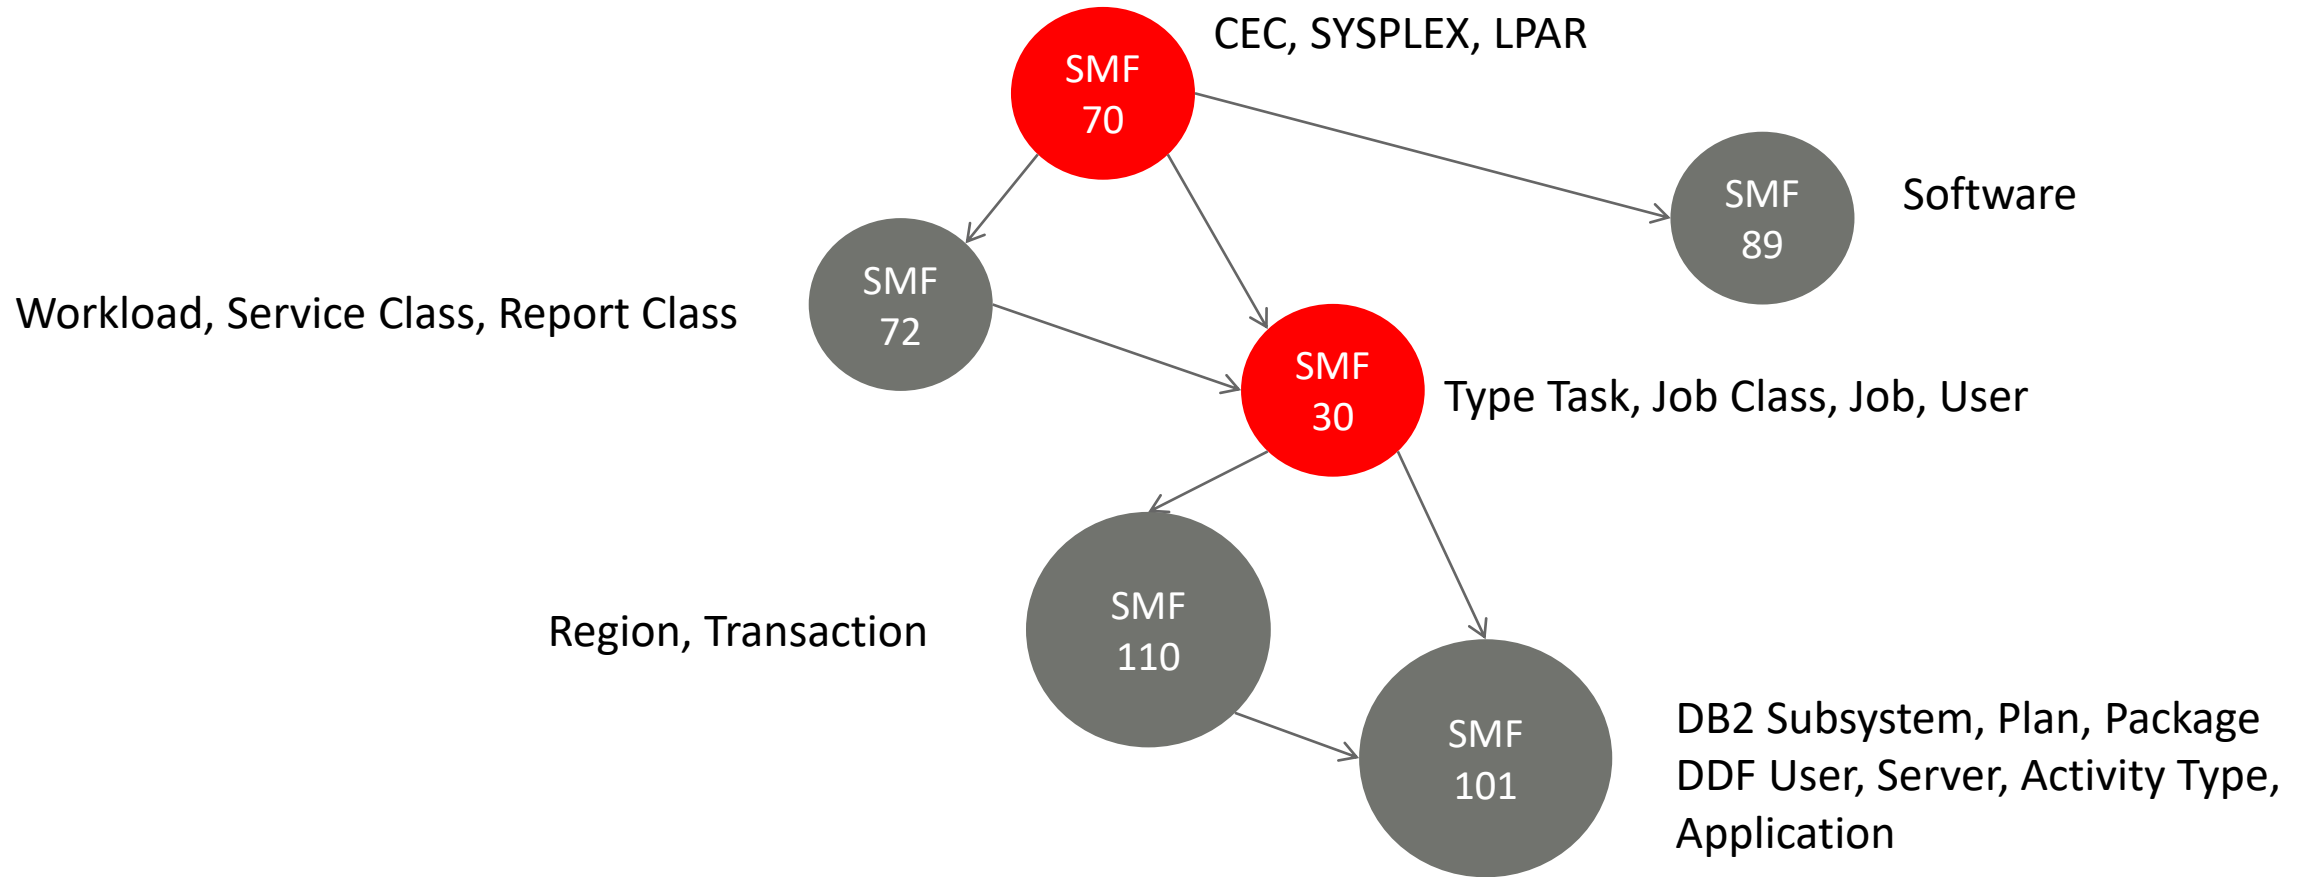
SMF 70 + 89





SMF 89 registers which software products are running on which LPARS. This is the basis for the SCRT report and IBM Software invoicing. The monthly charge is determined by the peak 4HRA MSU usage for all the LPARS running a given product.

Here we see the impact of a developer starting CICS on a test LPAR (red)

Example: Batch jobs requires tuning




SMF70



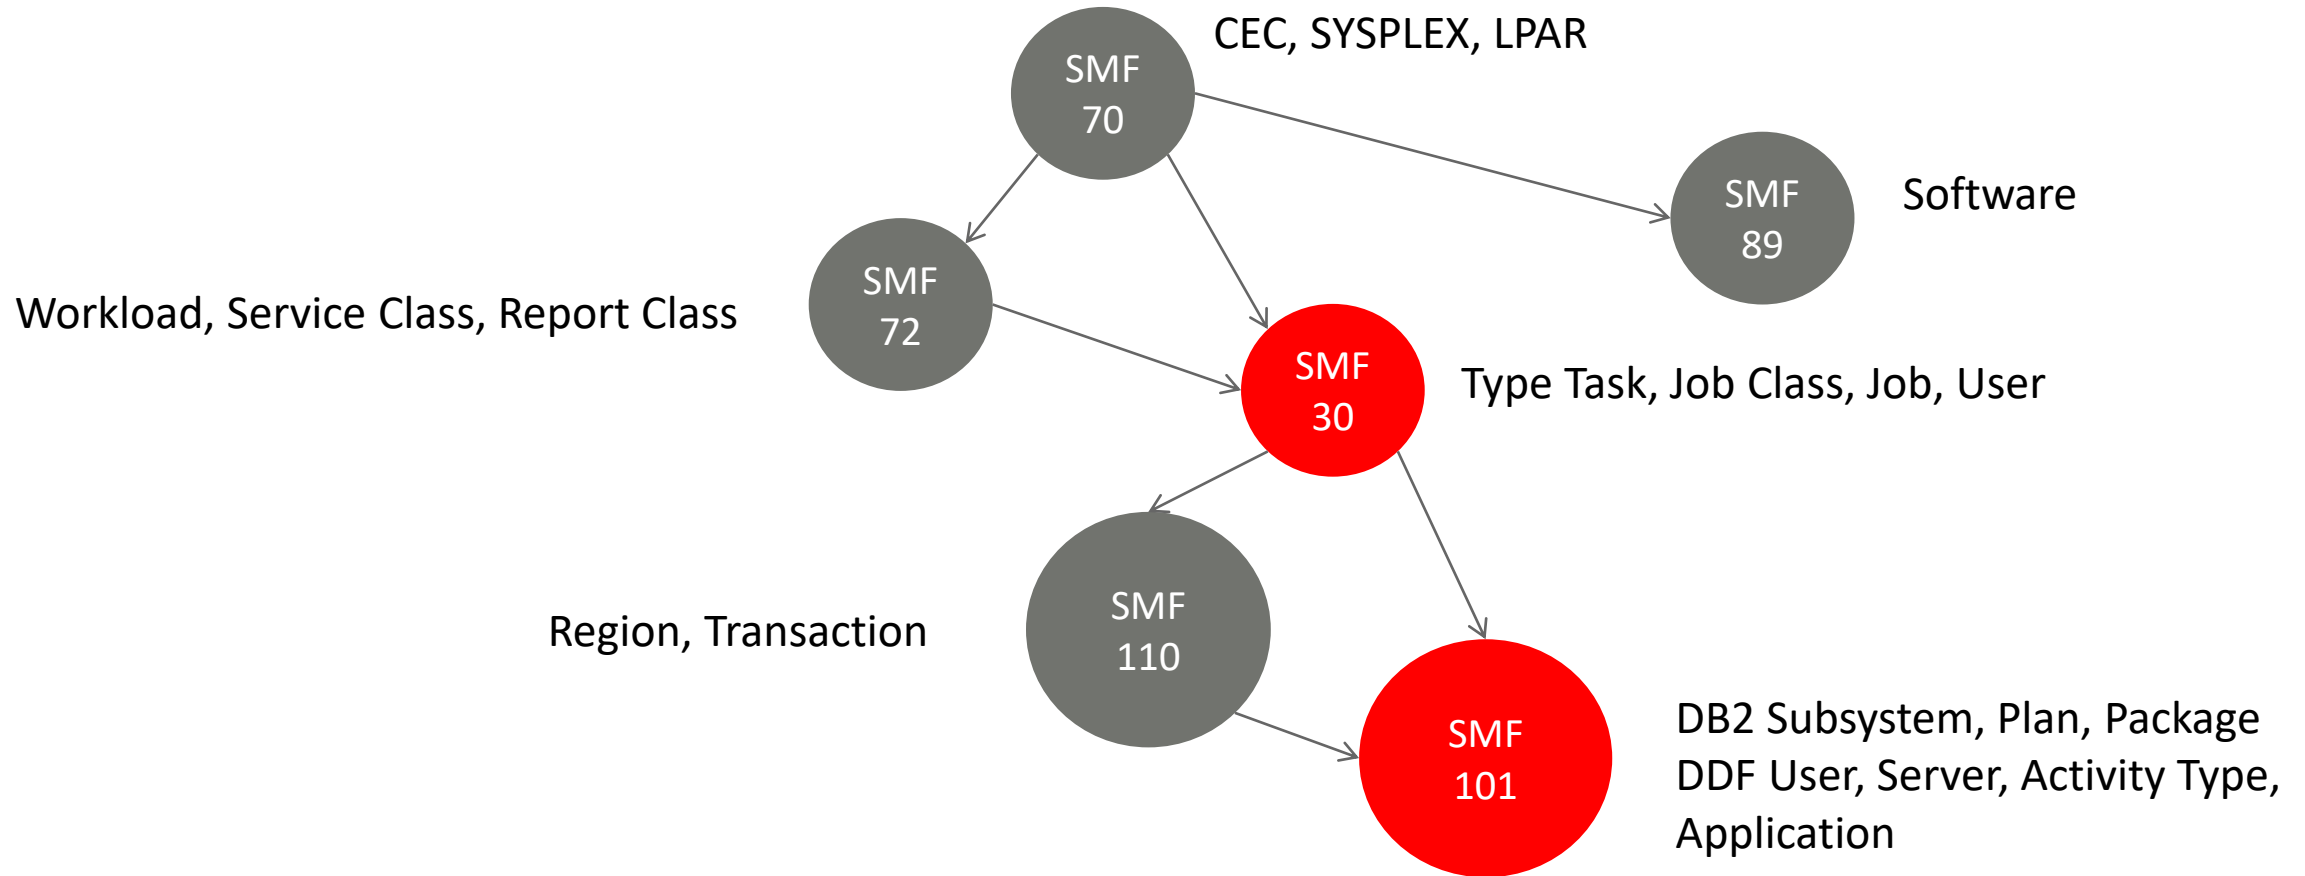
Even without the Dev and Test LPARS, September 19th is still the peak day.

SMF 30


Total	437.9	100
JOB12345	87.7	20
	36.7	8
	26.4	6
	20.5	5
	20.3	5
	15.3	3
	14.4	3
	13.0	3
	11.0	3
	10.8	2
	10.5	2
	10.2	2
	8.6	2
	8.5	2
	8.1	2



Example: DDF



SMF30




In SMF 30 we can see that DNSDIST (DDF) has a spike on October 26

SMF101

Hourly Package MIPS Usage

⏮ ⏭ 📊 ✕




In SMF 101 we see that the spike is primarily due to the SYSLH200 package. As we will see later, we can further classify the workload here based on server, user, application, transaction etc.

Understanding the impact of change

Understand the impact of changes

- Compare two points in time
- Look at trends over time
- Drill down to the details

ⓘ Average CPU-Seconds per selected CICS Transaction(s) over Time



ⓘ Growth in Top MIPS-consuming Trx

	—	2015	Growth	
	—	Apr		
		23	24	
M0UÆ	15.49	35.43	↑	129 %
M3CJ	19.73	21.83	→	11 %
M7CØ	7.06	9.02	→	28 %
M8UQ	4.38	3.27	↓	-25 %
MVTJ	3.78	2.72	↓	-28 %

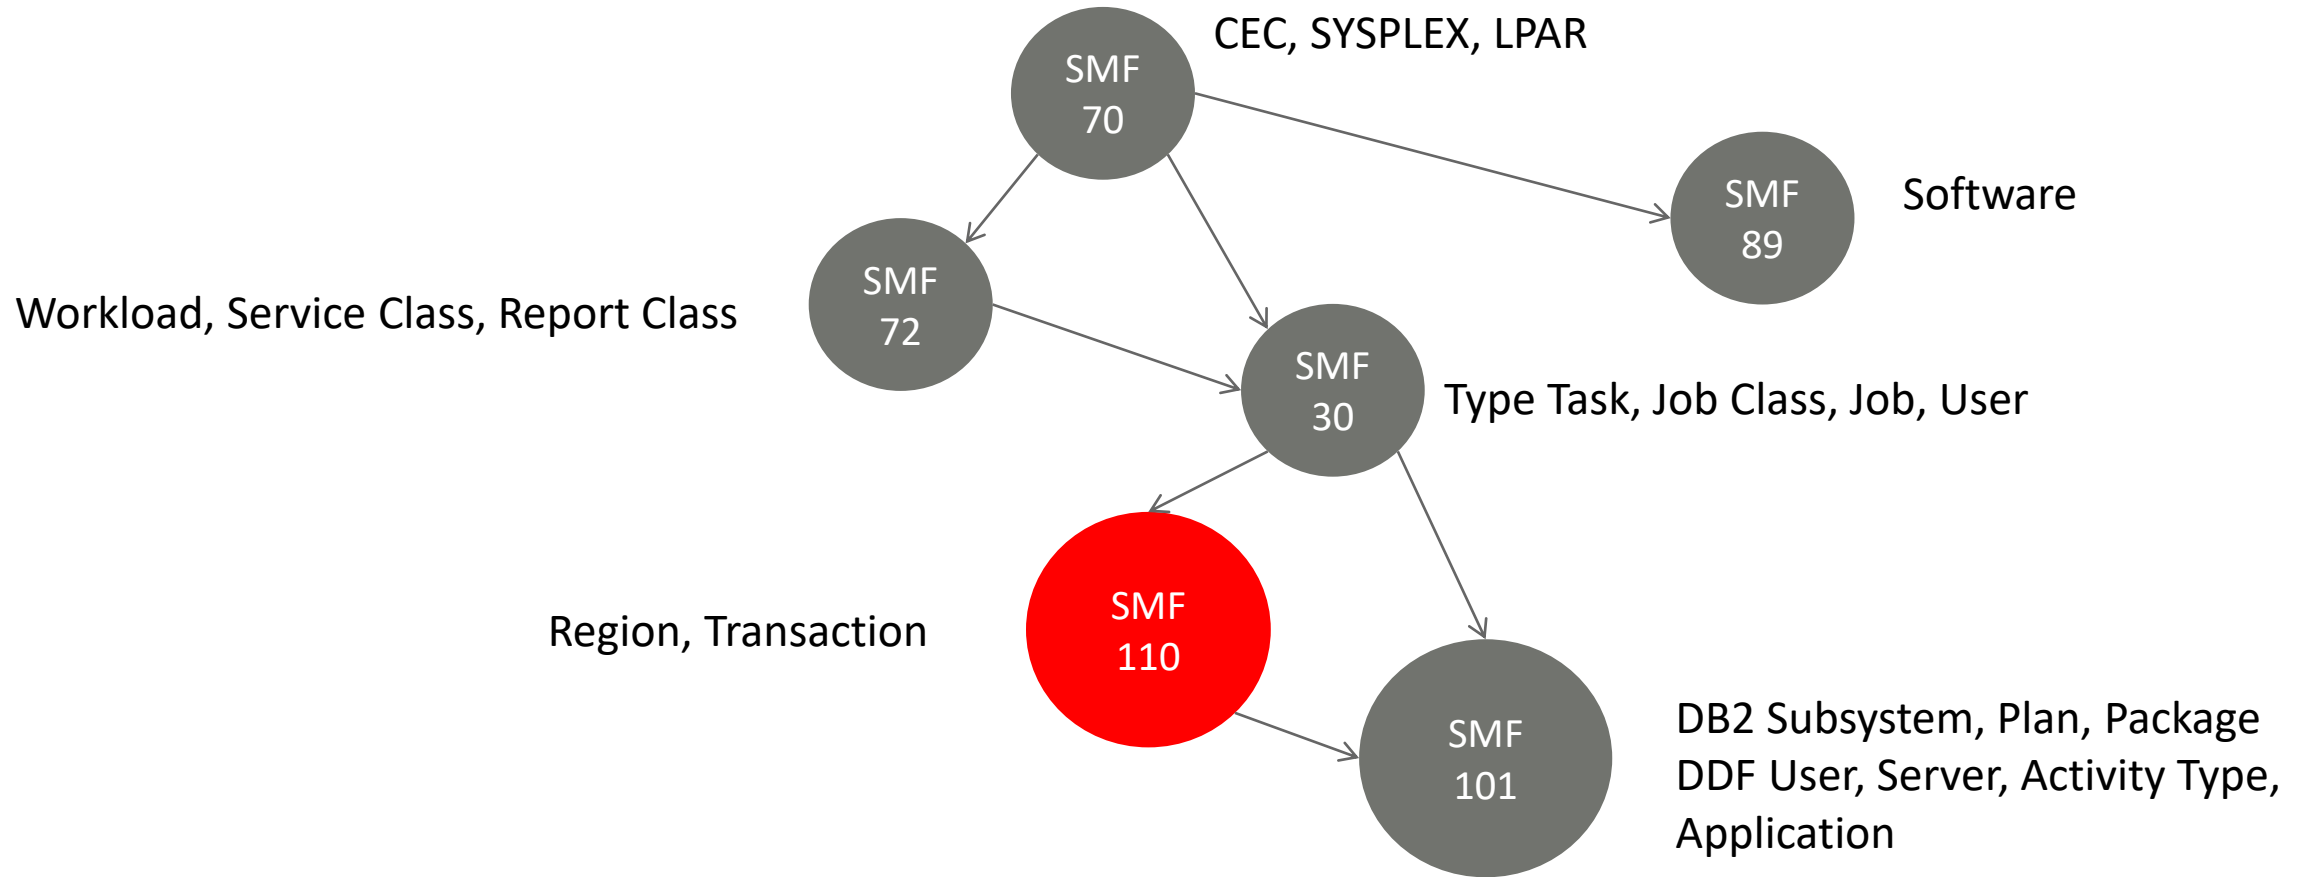
ⓘ Growth in # Trx

	—	2015	Growth	
	—	Apr		
		23	24	
M0UÆ	68,490	57,396	↓	-16 %

ⓘ Growth in heaviest Trx (CPUsec / Trx)

	—	2015	Growth	
	—	Apr		
		23	24	
M0UÆ	0.008	0.022	↑	173 %

Example: CICS



SMF 110

① Growth in Top MIPS-consuming Trx

	—	2015	Growth	
	—	Apr		
	23	24		
M0UÆ	15.49	35.43	↑ 129 %	^
M3CJ	19.73	21.83	→ 11 %	
M7CØ	7.06	9.02	→ 28 %	
M8UQ	4.38	3.27	↓ -25 %	
MVTJ	3.78	2.72	↓ -28 %	
M4CJ	2.89	3.15	→ 9 %	
VJ7C	.09	3.00	↑ 3,323 %	
M2Ø7	.82	.85	→ 3 %	
M9UØ	.53	.82	↑ 56 %	
VTAN	.55	.54	→ -2 %	
M3UJ	.55	.51	→ -8 %	
VH7E	.77	.23	↓ -70 %	
VL7E	.51	.40	↓ -21 %	

① Growth in # Trx

	—	2015	Growth	
	—	Apr		
	23	24		
M0UÆ	68,490	57,396	↓ -16 %	^
M3CJ	34,518	36,815	→ 7 %	
M4CJ	20,204	21,547	→ 7 %	
EZUL	6,183	4,946	↓ -20 %	
FYKE	6,183	4,946	↓ -20 %	
VL7E	2,594	2,036	↓ -22 %	
VI70	2,544	2,005	↓ -21 %	
VI7V	2,445	1,919	↓ -22 %	
M7CØ	1,745	1,597	→ -8 %	
VG71	1,834	1,499	↓ -18 %	
VG7O	1,802	1,479	↓ -18 %	
M6CQ	1,578	1,646	→ 4 %	
VH7M	1,624	1,318	↓ -19 %	

① Growth in heaviest Trx (CPUsec / Trx)

	—	2015	Growth	
	—	Apr		
	23	24		
VJ7C	0.045	2.209	↑ 4,860 %	^
M8UQ	0.520	0.346	↓ -34 %	
VR3R	0.338	0.346	→ 2 %	
VK7Q	0.328	0.237	↓ -28 %	
VH7E	0.254	0.164	↓ -35 %	
M3UJ	0.287	0.147	↓ -49 %	
MVTJ	0.225	0.166	↓ -26 %	
M7CØ	0.147	0.205	→ 40 %	
M2Ø7	0.157	0.162	→ 3 %	
VK7H	0.120	0.151	→ 26 %	
M9UØ	0.100	0.157	↑ 57 %	
M0UØ	0.109	0.085	↓ -22 %	
U282	0.024	0.114	↑ 378 %	

① Growth in Trx Response Time (Seconds)

	—	2015	Growth	
	—	Apr		
	23	24		
EQQJ	1,887.438	1,887.439	→ 0 %	^
M3UJ	467.410	262.244	↓ -44 %	
M8UQ	113.041	96.748	→ -14 %	
MVTJ	48.334	45.154	→ -7 %	
HIGU	44.685	45.158	→ 1 %	
M6UJ	37.945	35.984	→ -5 %	
M5UQ	36.615	33.343	→ -9 %	
MVKJ	34.361	33.258	→ -3 %	
U284	34.191	32.132	→ -6 %	
U286	42.226	30.022	↓ -29 %	
M7UJ	31.799	31.429	→ -1 %	
M9UØ	30.839	31.761	→ 3 %	
MXTJ	30.977	30.384	→ -2 %	

Using the SMF 110 we can compare two points in time. Here we see change in terms of total MIPS, number of transactions, MIPS per transaction, and response time.

SMF 110

① Growth in Top MIPS-consuming Trx

	—	2015		Growth
	—	Apr		
	23	24		
M0UÆ	15.49	35.43	↑	129 %
M3CJ	19.73	21.83	→	11 %
M7CØ	7.06	9.02	→	28 %
M8UQ	4.38	3.27	↓	-25 %
MVTJ	3.78	2.72	↓	-28 %
M4CJ	2.89	3.15	→	9 %
VJ7C	.09	3.00	↑	3,323 %

① Growth in # Trx

	—	2015		Growth
	—	Apr		
	23	24		
M0UÆ	68,490	57,396	↓	-16 %

① Growth in heaviest Trx (CPUsec / Trx)


	—	2015		Growth
	—	Apr		
	23	24		
M0UÆ	0.008	0.022	↑	173 %

① Growth in Trx Response Time (Seconds)

	—	2015		Growth
	—	Apr		
	23	24		
M0UÆ	0.121	0.137	→	13 %

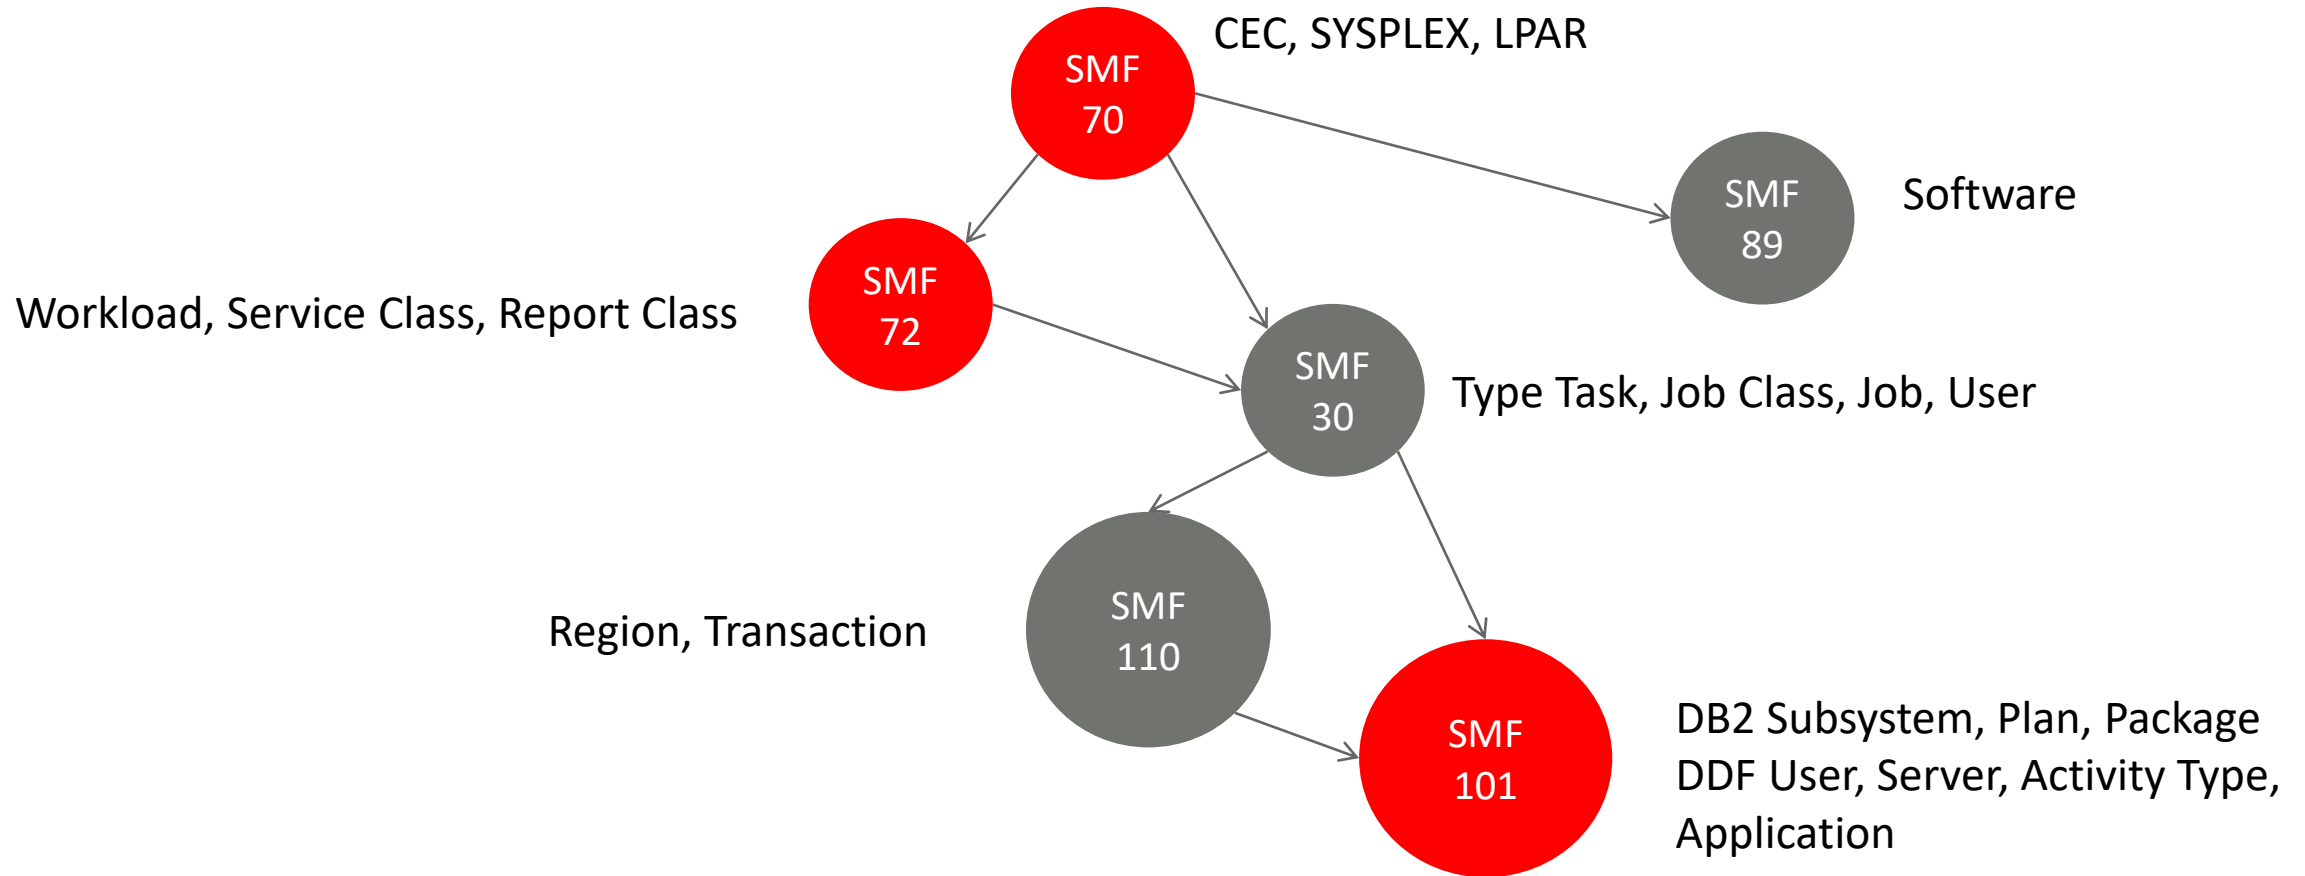
Selecting the M0UÆ transaction, we see that the MIPS increase is not due to increased transaction volumes, but rather due to a significant increase in the MIPS cost per transaction. Note that the response time was not significantly affected by this.

SMF 110




Looking at the MIPS per transaction for this transaction we see a significant change between April 23 and 24, probably indicating an application change.

Example: DDF




SMF 70




Analysis of SMF70 shows that the average MIPS usage during working-hours increasing month by month

SMF 72




From SMF72 the primary cause of this increase is seen to be the DDF Workload that has increased from 640 to over 1000 MIPS during the 9 month period.

SMF 101



SMF 101 allows us to analyse the change during the period from a number of angles. Here by subsystem where we see that the change is isolated to the DB2A subsystem


SMF 101



Drilling down into the DB2A subsystem by package, we see that SYSTAT (stored procedures) accounts for the largest absolute increase, but package1 and especially package4 have large percentage increase - indicating a probable application change

SMF 101

Mips per End User Transaction by calendar Date



Further analysis by End User Transaction gives additional insight. Note in this example we only see 'standard' transactions such as the IBM Data Server Driver for JDBC and SQLJ. Note that the application developers have the ability to overwrite this with other information using the **setClientProgramName** method.

SMF 101


Other areas of interest in the SMF 101 regarding DDF include:

- Server (ip address)
 - Mapping server ip addresses to business areas or applications can be useful
 - But note this may be a gateway server and may not say much about where the workload is coming from
- Userid or Authorization id
 - But note that the application may be running under a standard 'production' user
 - Running different applications under different 'production' userids can be an easy way to further split out the workload coming from DDF
 - Mapping userids to business areas or applications can also be useful here

Tune the infrastructure

Tune Infrastructure


- Understand what is contributing to peak utilization
- Tune system parameters to reduce consumption during peak
- Move things away from the peak:
 - Manual capping
 - Automatic (dynamic) capping
 - Rescheduling



Explain the costs to the business

Explain the capacity costs to the business

- Map the technical measurements to something the business or developers can understand:
 - Application
 - Organization
 - Environment
 - Service Provider
- Translate the technical measurements such as MIPS or MSU into money



Conclusions

- SMF and other capacity data can form the basis of a fruitful dialog with development
- Making the data available in terms that are understandable to the developers helps
- Developers can help make the data more useful through business mapping and by putting additional information into the SMF
- Making self service reporting available to the developers helps
- Internal charge back or show back can be built on this data to increase the level of attention and motivation from development

We want your feedback!

- Please submit your feedback online at
 - <http://conferences.gse.org.uk/2018/feedback/LM>
- Paper feedback forms are also available from the Chair person
- This session is LM

