Available on CMS information server **CMS IN -2009/030**

CMS Internal Note

The content of this note is intended for CMS internal use and distribution only

20 August 2009 (v3, 21 August 2009)

Status of Data- and Work-Flow Planning for the CMS SUSY Group

C. Auterman, A. Bhatti, O. Buchmueller, M. Chiorboli, R. Demina,
M. Pierini, J. Richman, F. Ronga, M. Spiropulu, D. Stuart, A. Tapper
(on behalf of the SUSY group)

Abstract

This note presents a first version of a possible data- and work-flow strategy for physics analyses in the CMS SUSY group. In the context of a broad set of searches based on topological SUSY signatures, we discuss requirements for secondary datasets and group skims. We also describe a possible strategy for data reduction using a group-specific PAT-tuple production, which is tailored especially to the early commissioning phase of the SUSY analyses.

CMS Internal Note

The content of this note is intended for CMS internal use and distribution only

20 August 2009

Status of Data- and Work-Flow Planning for the CMS SUSY Group

C. Auterman[a], A. Bhatti[b], O. Buchmueller[c], M. Chiorboli[d], R. Demina[e], M. Pierini[f], J. Richman[g], F. Ronga[h], M. Spiropulu[d,f], D. Stuart[g], A. Tapper[c] – on behalf of the SUSY group.

a-Hamburg University, b-Rockefeller University, c-Imperial College London, d-California Institute of Technology, e-University of Rochester, f-CERN, g-University of California, Santa Barbara, h-Institute for Particle Physics, ETH Zürich

Abstract. This note presents a first version of a possible data- and work-flow strategy for physics analyses in the CMS SUSY group. In the context of a broad set of searches based on topological SUSY signatures, we discuss requirements for secondary datasets and group skims. We also describe a possible strategy for data reduction using a group-specific PAT-tuple production, which is tailored especially to the early commissioning phase of the SUSY analyses.

1. Introduction

This document describes a set of triggers, datasets, and skims that could support the study of a broad range of SUSY signatures in the first LHC physics run. We also analyse a number of practical considerations related to the organisation and handling of these data samples.

We begin in Section 2 with a discussion of the expected overall data-flow in CMS, which is based on primary datasets, secondary datasets, and central skims. We discuss the implications of this structure for SUSY analyses. Section 3 describes the requirements for the SUSY Reference Analyses (RAs), which form a coherent set of SUSY searches for the first run. For each RA, we list a set of triggers, data samples, and skims that would support the analysis. We present estimates for the disk space needed for each RA, along with the underlying assumptions for these estimates.

Section 4 presents a first picture of how the SUSY group skim production might be structured and discusses issues such as central vs. local production of samples, file transfers, and the evolution of the plan with increasing luminosity. Section 5 discusses our experience with producing Monte Carlo samples in a particular data format, the PAT-tuple. This format could be used for the SUSY group skims, but other formats, such as the AOD, could also be suitable. Finally, we present in Section 6 an overall summary of the SUSY group's skim strategy and a list of questions to be addressed in the future.

We emphasise that this document is intended to be a contribution to an overall, CMS-wide discussion about how to manage most effectively the samples used by the Physics Analysis Groups. The SUSY group supports a common plan that would facilitate the sharing of resources across the different groups.

2. Expected CMS general data flow

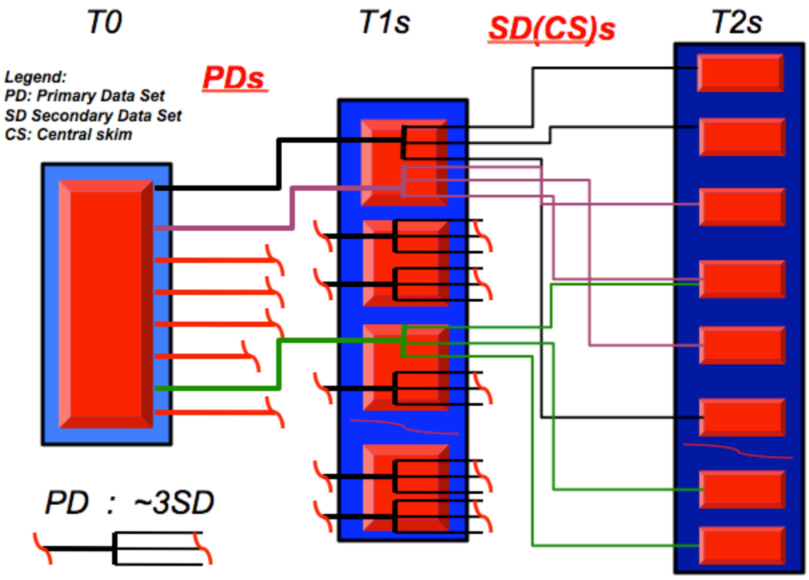

Figure 1. Schematic illustration of the general dataflow in CMS from T0 to T2. Approximately 10 PDs are reconstructed at the T0 and then distributed to the T1s. At the T1s, the PDs are split into a total of 30 to 40 SDs (CSs), which are then distributed to T2s for physics analysis.

Figure 1 illustrates the expected offline data flow of CMS. Primary Datasets (PDs) in CMS are reconstructed at the T0 at CERN and then distributed to the different CMS T1s. PDs are event samples defined according to groups of physics trigger bits, and the exact specification of the grouping is driven mainly by the technical constraints of the T0 reconstruction (O(10) PDs of approximately equal size). In addition, the PD definitions might change with increasing luminosity. Therefore, they are not expected to be optimal for physics analysis.

A more suitable concept for physics analysis in the early days is Secondary Datasets (SDs). These event samples are produced at T1s by splitting a PD into smaller event samples defined by a subset of the original trigger bits or even by a single trigger. Like PDs, SDs are based solely on trigger information, so it is possible to consistently re-reconstruct a single SD. In addition to SDs, there also is a possibility to produce Central Skims (CSs) at the T1 level. Like SDs, CSs would be based on a single PD but would carry out an event selection based on reconstructed offline quantities or quantities in the trigger objects, not just trigger bits. Therefore, the event content of a CS may not be invariant under re-reconstruction, whereas SDs are invariant. Both SDs and CSs are produced at the T1s and then distributed to T2s.

During the early phase of data taking, it is foreseen to split a PD into not more than three SDs (or CSs). This number is expected to increase in a later stage of the running. Furthermore, as long as enough disk space resources are available, SDs and CSs will be distributed in RECO format (with roughly 400 kB/event) to the T2s. After this point, a significant reduction in data size must be achieved, and it is likely that some or even all of the data distribution to T2s will be performed using an analysis object with a size of around 100 kB/event.

3. Trigger and Dataset Requirements for SUSY Reference Analyses

3.1. SUSY Reference Analyses

The SUSY group has formulated a plan to search coherently for a broad range of generic, topological signatures. These signatures span essentially all important physics objects and include multijets, photons, single leptons, dileptons, and trileptons, each of which is produced in association with missing transverse energy (MET). To organise this wide variety of topologies, the SUSY group has defined a concept of reference analyses (RAs), which are described in detail elsewhere (see, for example, [1]). Searches for SUSY in hadronic final states include:

RA1: Exclusive n-jet analysis

RA2: Inclusive ≥3 jet analysis

RA3: Diphoton+jets+MET, Photon+jets+MET

Searches involving leptons in the final state are grouped in four reference analyses:

RA4: Single lepton+jets analysis

RA5: Same sign dilepton analysis

RA6: Opposite sign dilepton analysis

RA7: Trilepton analysis

Each RA can include multiple signatures due to, for example, different lepton flavors, b-jet or τ tagging, or top reconstruction. In the following, we discuss the requirements for triggers and analysis datasets that arise from the RA strategy.

3.2. SUSY group resources

The constraints imposed by the computing resources available to physics groups must be taken into account when defining a sound data flow for physics analysis.

Currently, the SUSY group has been assigned 30 TB of disk at five T2 sites: Florida, London, Aachen, Vienna, and Bari, for group usage. Thus, the amount of disk space for storage of group skims is approximately 150 TB. However, it is expected that in the forthcoming refinement of the T2 group quota assignment, the different physics groups will get additional disk resources for the start of data taking. Therefore, 150 TB of overall SUSY T2 space represents a conservative estimate of the available resources at start-up. This storage corresponds to about 375 M events in RECO format or 1.5 G events in a 100 kB/event format.

3.3. Trigger Path Performance Studies on SUSY benchmarks

In order to identify the most efficient triggers for SUSY searches, the performance of the HLT trigger paths in the start-up trigger table [2] was studied for a set of mSUGRA models, as well as for less constrained SUSY scenarios. The efficiency for each trigger path was determined for each of the different reference analysis selections as defined in the last column of Table 2. For reasons of simplicity and robustness, the SUSY trigger strategy will mainly focus on single trigger bits at start-up.

Figure 2 presents the signal efficiencies for an RA1-type all-hadronic signal selection for a variety of triggers listed on the horizontal axis and for a selection of SUSY signal scenarios represented by different colours. The hadronic (Jet/MET/HT) trigger paths are the highest efficiency ones and in many cases the HLT_Jet_110 and HLT_HT_200 paths are optimal.

Figures 3 and 4 present the results for the leptonic RA4-type of signal selection in the electron and muon channels respectively. We note that in many cases the hadronic trigger paths are also very efficient for the RA4 analyses with an electron or a muon required in the final state. While for most models studied the single lepton triggers perform best, a cross-trigger with a lower p_T lepton and an HT requirement combined in a logical or with the nominal single lepton triggers could increase the signal efficiency by 15% to 30% without adding large overall rate and disk space requirements. For example, for the LM9 signal scenario the HLT_Ele10_HT180 path could be used to improve the trigger efficiency at low lepton p_T compared to the HLT_Ele15 path. This emphasises the need for a variety of trigger paths for SUSY searches.

In addition, the overlap of signal events in the different trigger paths of the start-up HLT table was studied targeting an optimal strategy for increasing the trigger efficiency for the signal. As an example, the overlap for the LM9 signal scenario between HLT_Ele15_LooseTrackIso and HLT_Ele10_HT180 is 76%; hence it is possible to gain up to 24% by also including the cross-trigger with the nominal single lepton trigger. Similarly, the overlap between HLT_Ele20 and HLT_Ele10_HT180 is 83%, so it is again possible to gain 17% efficiency with respect to the single electron trigger.

The results of these studies confirm that the SUSY group can rely at start-up on the simple and robust single-object triggers. However, as suggested by some of the illustrative examples discussed in this section, additional signal efficiency might be gained by combining single-object triggers with cross-object triggers. The use of such trigger combinations to complement the single-object trigger strategy for early SUSY searches is currently under investigation and therefore the consequences for skims and datasets must be considered. However, for the purpose of this note it is enough to conclude that single-object triggers are sufficient for the commissioning and execution of the early SUSY searches. A detailed discussion of the full trigger strategy, including efficiency measurements is beyond the scope of this document and will be outlined in more detail in an internal note that is in preparation.

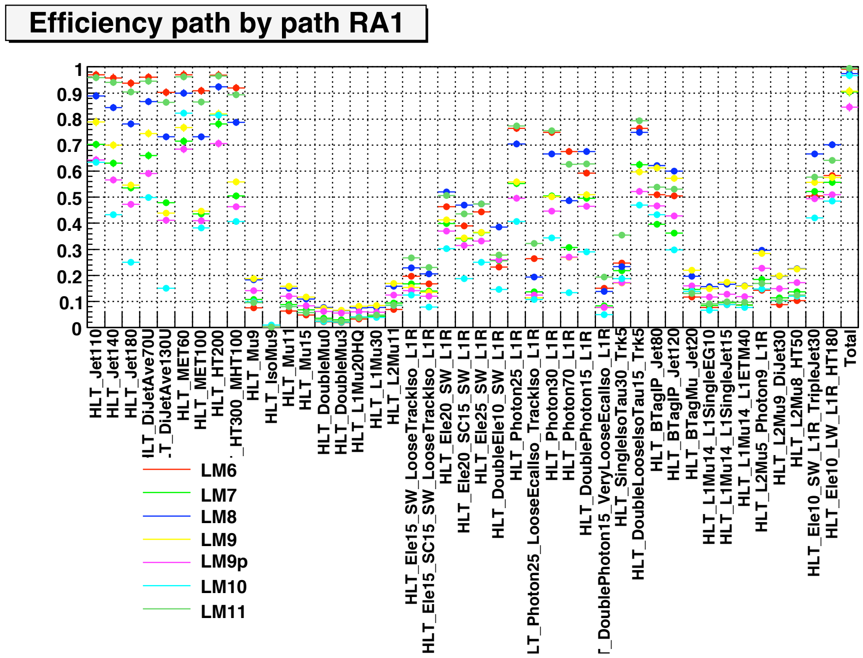

Figure 2. Trigger efficiency for all-hadronic (RA1) analysis for a set of mSUGRA benchmark points. The x-axis lists the different triggers defined in [2], while colours mark the different SUSY scenarios. As expected, the hadronic (JET/MET/HT) paths have the highest efficiency, with the HLT_JET_110 and HLT_HT200 being optimal in many cases.

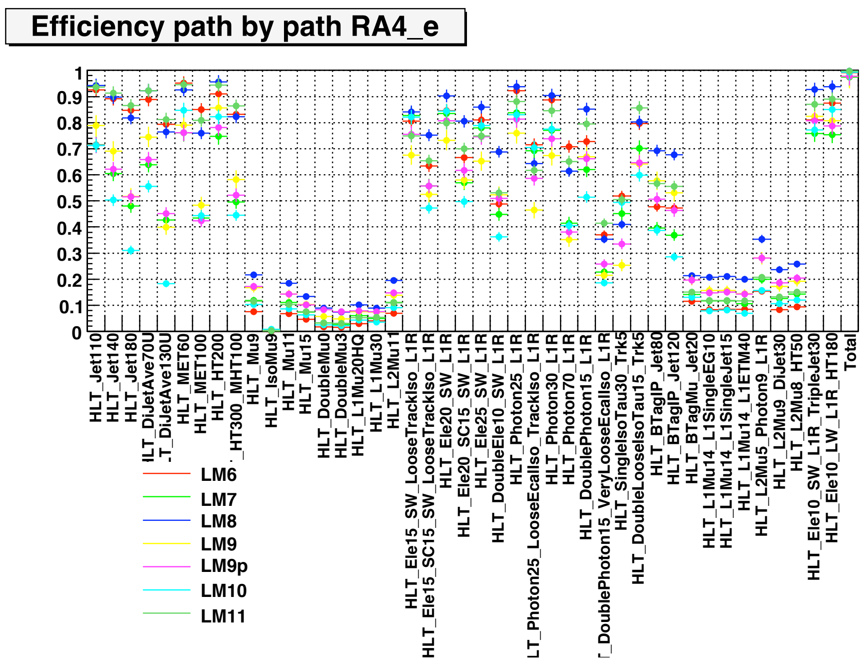

Figure 3. Performance of trigger paths for electron final states (RA4 analyses) and for a set of mSUGRA benchmarks points.

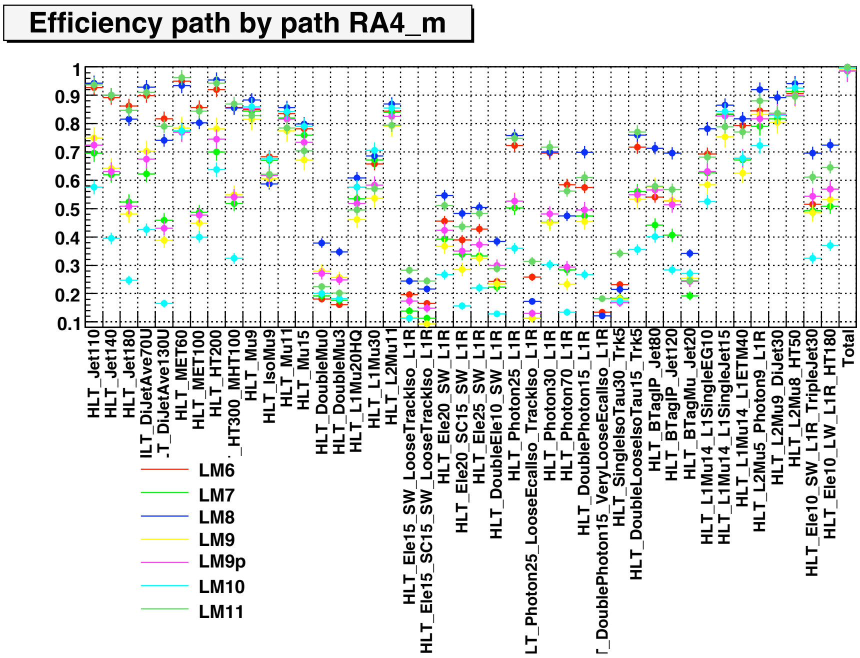

Figure 4. Performance of trigger paths for muon final states (RA4 analyses) and for a set of mSUGRA benchmarks points.

3.4. Trigger and Dataset Requirements

Especially at the beginning of data taking, but possibly also at a later stage, it will be important for the SUSY group to have access to all events taken by relevant SUSY triggers. For that reason, SDs will be the primary input source for SUSY group skims in the early days.

To provide input to the CMS-wide definition of the SDs, the SUSY group has studied in detail the set of triggers and associated disk space required for early SUSY searches. The results of these studies are summarised in Table 1. It lists the important triggers for each of the SUSY RAs as well as the corresponding amount of disk space required to store data in the 400 kB RECO and in the 80 kB PAT-tuple (or equivalent) format. The estimates listed in Table 1 are based on the following assumptions[1]:

- The instantaneous luminosity is $L_{ins} = 1E31\ cm^{-2}\ sec^{-1}$
- Total integrated luminosity is $L_t = 100\ pb^{-1}$
- HLT trigger rates (*R, Hz*) are taken from [2]
- PAT-tuple event size is $S_{PAT\text{-}tuple} = 80\ kB$ (see Section 5)
- RECO event size is $S_{RECO} = 400\ kB$

[1] Total disk space needed for each data set (*D*) is the number of events times the size of each event: $D=N_{ev}S_{ev}$. The number of events is equal to the rate of this trigger times the effective time of data collection, or the ratio of integrated luminosity to instantaneous one:
$N_{ev}=RL_t/L_{inst}=R(100\ pb^{-1}/10^{31}\ cm^{-2}sec^{-1}) = R(10^{38}\ cm^{-2}/10^{31}\ cm^{-2}sec^{-1})=R(Hz)10^7\ sec$
Thus, the disk space needed for storage of PAT-tuple events is $D_{PAT\text{-}tuple}=0.8(TB*sec)R(Hz)$,
and for RECO events it is $D_{RECO}=4(TB*sec)R(Hz)$.

Table 1. Triggers for SUSY Reference Analyses (RAs). The triggers listed are taken from the current 1E31 trigger menu and should be understood as examples of triggers needed for key SUSY searches. The number of events and disk space estimates are evaluated for 100 pb^{-1} of integrated luminosity. The prefix HLT_ has been removed from these trigger names.

SUSY RA	Triggers	Rate [Hz]	# of events (million)	RECO size [TB]	PAT-tuple size [TB]
1,2	Jet_110	7.1	71	28.4	5.7
	MET_50	0.5	5	2	0.4
	HT_300_MHT_100	0.1	1	0.4	0.1
	Total hadronic triggers	**7.7**	**77**	**30.8**	**6.2**
3	Photon25_L1R	20.5	205	82	16.4
	DoublePhoton15_L1R	4.7	47	18.8	3.8
	Total photon triggers	**25.2**	**252**	**100.8**	**20.2**
4,5,6,7	Mu9	12	120	48	9.6
	Double_Mu3	2.3	23	9.2	1.8
	Total muon triggers	**14.3**	**143**	**57.2**	**11.4**
	Ele15_SW_LooseTrackIso_L1R	14.6	146	58.4	11.7
	DoubleEle10_SW_L1R	2	20	8	1.6
	Total electron triggers	**16.6**	**166**	**66.4**	**13.3**
	Total SUSY triggers	**63.8**	**638**	**255.2**	**51.1**

The triggers listed in Table 1 are from the current 1E31 menu and should be understood as examples of important triggers that are required for SUSY analysis commissioning and execution. We note that HLT_Jet_110 is the lowest threshold un-prescaled single-jet trigger in 1E31 menu. Its rate, about 7 Hz, dominates the estimated rate of the triggers proposed for RA1 and RA2. In the current menu there is also a HLT_HT_200 trigger that has a rate of 16 Hz. Its rate is mainly dominated by QCD dijet events and therefore this total hadronic jet energy trigger is highly correlated with the single-jet trigger. Both triggers are very efficient for SUSY searches and due to their high overlap it is assumed that these two triggers will be merged into a single SD. Table 1 lists only the HLT_Jet_110 trigger that has the lower rate and therefore is more robust. In the case that these triggers are merged into a single SD, the amount of disk space required would be approximately 60 TB instead of the currently listed 28 TB for HLT_Jet_110 only.

The additional triggers, HLT_HT_300_MHT_100 and HLT_MET_50, will enable us to study high MET/MHT events, but the behavior of these triggers is more complex than the single-jet trigger.

RA3 will make use of the lowest threshold, un-prescaled single-photon trigger that is practical, currently thought to be HLT_Photon25_L1R. The rate for this trigger is substantial, around 20 Hz. Therefore, the total photon triggers listed in Table 1 require approximately a factor of two more disk space than the corresponding electron and muon triggers. When disk space resources become an issue, it might be necessary to revert to a higher threshold of this trigger possibly even in combination with the di-photon trigger. In addition, further refinements of the trigger menu might lead to a more democratic trigger rate allocation across the single-object triggers.

RA4, the single lepton SUSY analysis, requires both muon and electron triggers. The single muon trigger, HLT_Mu9, and the single electron trigger, HLT_Ele15_SW_LooseTrackIso_L1R have comparable rates, 12 Hz and 15 Hz, respectively. While adjusting the muon energy threshold can efficiently control the rate of the single muon trigger, ensuring robust triggering at low electron energies is more difficult. In case the rate of the single electron trigger becomes too high, it is possible to add an additional HT requirement to control the rate by keeping the same or even a lower electron energy threshold. An example for such a cross object trigger in the current 1E31 menu is HLT_Ele10_SW_L1R_HT180 with a trigger rate of only 10 Hz at a electron energy threshold of 10 GeV.

For RA4, RA5, and RA6, double-muon and double-electron triggers will be used in addition to the single-lepton triggers, and are expected to have fairly low rates, around 2 Hz each.

While for reasons of simplicity and robustness the SUSY trigger strategy will mainly focus on single trigger bits at start-up, it seems possible that the simple lepton triggers can be combined with cross-triggers without much additional disk space required, potentially boosting the signal efficiency for leptonic analysis as discussed for a few examples in Sections 3.3. The use of simple multi-trigger paths for some of the RAs is currently under investigation.

The total rate associated with the triggers needed for SUSY analyses is around 72 Hz, resulting in a sample of around 640 M events in a 100 pb^{-1} run. The amount of disk space needed to store all events relevant for SUSY analyses in RECO format is about 250 TB, while the 100 kB analysis object samples would require about 50 TB. The SUSY T2 nominal storage of 150 TB can accommodate the 100 kB analysis object samples.

The disk space estimates listed in Table 1 assume that each single trigger will be made available as a single SD. Therefore, the event overlap of these triggers propagates directly into the overall estimate. If, however, some of the required SUSY triggers are grouped into a single SD, due to the non-zero overlap of these triggers, the required disk resources for this SD will be smaller than the individual sum. In this respect, the estimates in Table 1 represent conservative estimates of the overall disk resources needed to store events from these triggers at T2s.

Furthermore, it is important to note that all of the SUSY triggers are also highly relevant for other CMS physics analyses. Therefore, the SUSY group assumes that the SDs containing these triggers will count against the central CMS quota when distributed to T2s.

In summary, the SUSY group requests the following triggers to be made available in the form of SDs:

- Single Jet Trigger (plus higher threshold backups)
- MET Trigger[2] (plus higher threshold backups)
- Single Muon Trigger (plus higher threshold backups)
- Double Muon Trigger (plus higher threshold backups)
- Single Electron Trigger (plus higher threshold backups)
- Double Electron Trigger (plus higher threshold backups)
- Single Photon Trigger (plus higher threshold backups)
- Double Photon Trigger (plus higher threshold backups)

Here the generic name represents a trigger with a given threshold. This threshold can change with increasing instantaneous luminosity or changing machine/detector conditions and therefore the generic acronym should be understood as the trigger with the lowest, un-prescaled threshold that is available in a given SD. The backup triggers are therefore un-prescaled thresholds at higher values of the same trigger.

[2] Trigger paths based on MHT are also in development and will be studied when available.

4. Data and workflow for SUSY group skim production

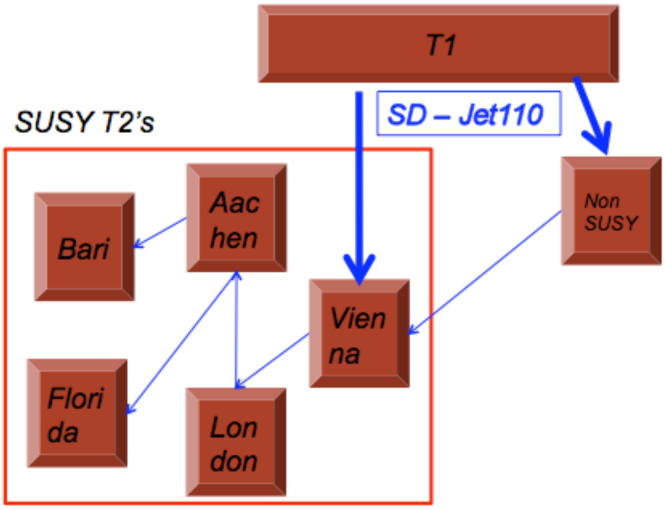

Figure 5. Schematic illustration of the possible data flow for group skim production. A secondary dataset (here containing the Jet110 trigger) is distributed to a T2. In case the SD is located at a SUSY T2, groups skims derived from the SD are stored directly at that T2. If the SD is located at a T2 not belonging to the SUSY analysis centres, the group skim is produced at that non-SUSY T2 and staged out directly to a SUSY T2. After the production of the SUSY group skim, the skim data are replicated via T2-T2 transfer to other SUSY T2s.

To facilitate the data access and to meet the constraints of the finite disk space resources of the group, it is important to reduce the size of the RECO format (400 kB/event) to a 50 to 100 kB analysis object format. This format could be the centrally produced AOD. In the case that there is no central production of the AOD, or if the AOD does not contain the information needed for all important SUSY analyses, the SUSY group would like to possess the flexibility to produce its own 100 kB analysis object that is tailored explicitly to the needs of the early searches. The production of this analysis object would be a group organised effort, integrated directly into the SUSY group skim activity.

Figure 5 shows a schematic illustration of the data- and work-flow for the SUSY group skim production, which can be summarised in three steps:

1. Data transfer of a SUSY-relevant SD (e.g., Jet110) to a T2. The SD will count against the central CMS quota.
2. Production of the corresponding group skim, based on a 100 kB analysis object, at the T2. In case the T2 is a SUSY analysis centre, the group skim will be written directly to the local SUSY group space at this T2. If, however, the T2 hosting the SD is not a SUSY T2, the group skim will be produced at this T2 and staged out directly to group space located at one of the SUSY T2s.
3. Once stored in the group space, the group skim of a given SD will be distributed via T2-T2 (or T2-T1-T2) transfer. Depending on the importance of the group skim, several copies of it might be made available on different SUSY T2s. However, for backup reasons, there always will be at least two different SUSY T2s that will host the same group skim.

In the past few months, the SUSY group has exercised the above described group skim production. The chosen 100 kB analysis object format was a PAT-tuple format. The format and lessons learned during the group skim exercise are addressed in Section 5. It is important to stress that all of the skim strategy discussed in this section only makes an assumption on the size of the 100 kB analysis object. The exact choice of data format for this object is not relevant for this discussion.

4.1. SUSY group skim data management

As listed in Table 1, for an integrated luminosity of 100 pb^{-1} a full set of SUSY group skims requires approximately 50 TB of disk space. Each group skim is assumed to be available at two different SUSY T2s doubling the required resources. Furthermore, the SUSY group would like to have the possibility to host two different reconstruction versions of the same skim. Therefore, at least four different versions of a given group skim will be stored at the same time in SUSY group space. In addition, it is foreseen to have at least the equivalent of one full copy of a group skim in Monte Carlo data permanently available on disk.

For that reasons, the SUSY group space must host five to six versions of a group skim at any point in time. Assuming 150 TB of total group space, six versions of the set of SUSY group skims in Table1 listed can be accommodated up to an integrated luminosity of around 50 pb^{-1}.

In the course of the soon planned group resource adjustment it is likely that the SUSY group space quota will increase significantly. However, even by doubling the group space resource to 300 TB, it is unavoidable that somewhere around 50-100 pb^{-1} of integrated luminosity a refinement of the group skim strategy must take place. Besides data cutback through event content reduction (RECO to 100 kB analysis object), also event rejection need to become part of the group skim strategy.

Therefore, for higher luminosities it is foreseen to apply a minimal set of cuts in order to reduce the size of the group skims so they can be stored using the available disk resources. These cuts will evolve with the understanding of the physics analysis, the size of the available data, and the availability of disk space. An example of these cuts is given in Table 2. These cuts are very safe compared to proposed analyses (last column in Table 2). We have not yet evaluated the reduction in data size due to these cuts.

Table 2. Example of important SUSY triggers included in SDs, possible additional cuts for group skims when needed, current RA cuts for comparison. The trigger examples are taken from the current 1E31 menu and like the offline cuts for skims and RAs are subject to change.

RA#	Triggers included in SDs (RECO)	Additional cuts for Group Skims when needed	Final selection (subject to change), user defined format
1	JET110, HT_200 HT300_MHT100, MET50	Two loose jets with $\|\eta_j\|<3.0$, $p_{Tj}>80GeV$	Leading jet with $\|\eta_j\|<2.0$, $p_{Tj}>100GeV$; $H_T>350GeV$, $\alpha_T>0.55$, $MHT_{ratio}<1.25$
2	Same as RA1	Two loose jets with $\|\eta_j\|<3.0$, $p_{Tj}>80GeV$	At least three jets with $\|\eta_j\|<2.5$, $p_{Tj}>180$, 150, 50 GeV MET>200 GeV, min[Dphi(MET, jet)]>0.3
3	Photon25_L1R DoublePhoton15_L1R +hadronic triggers	Diphoton skim common with QCD	
4	Mu9 Ele15_SW_EleId_L1R +hadronic triggers	Central electron skim ($p_T>15GeV$); Central muon skim ($p_T>9GeV$), non-iso;	Loose electron $p_T>20$ GeV, $\|\eta_j\|<2.5$, RelIso<0.1 Global tight muon $p_T>20$ GeV, $\|\eta_j\|<2.1$, RelIso<0.1
5+6	Double_Mu3, Double_Ele10_SWL1R, Artificial emu HLT ($p_T$mu>5GeV, $p_T$ele>10GeV) +single lepton triggers +hadronic triggers	Two muons $p_T$mu>3GeV; Two electrons $p_T$ele>10GeV; $p_T$mu>5GeV, $p_T$ele>10GeV remove same flavor, opp sign with M(ll)<10GeV Possibly for high stat samples add 3 *loose* jets with $\|\eta_j\|<3.0$, $p_{Tj}>30GeV$, or two jets with Dphi(jet1, jet2)<2.7	Robust electron $p_T>10$ GeV, $\|\eta_j\|<2.5$, RelIso<0.1 Global tight muon $p_T>10GeV$, $\|\eta\|<2.1$, RelIso<0.1 + jet $\|\eta_j\|<2.4$, $p_{Tj}>50$ GeV
7	Use RA5+6		

5. PAT-tuple production

Our recent experience in comparing event yields within SUSY analysis subgroups has highlighted the need for a centrally defined set of algorithms, parameters, and data samples. Beyond defining the algorithms and samples, a further way to ensure consistency across the group and potentially save time and effort is the production and storage of common files. In this case, we produced PAT-tuples derived from RECO samples and stored them for use by SUSY group members.

The definition and production of the files was performed in three stages:

1. The PAT-tuple format was defined: which algorithms to use, with what settings, and which collections to save in the output.
2. A set of samples to process then was selected and the production itself was performed. The resulting files were then distributed over the SUSY T2s.
3. In parallel, the files (produced during the two above stages) were validated and the procedures developed were evaluated.

These stages are described in more detail in the following sections. It should be noted that much of the experience gained in this effort is directly relevant to the issues that will be faced in executing group skims on collision data. Full details can be found at [3].

5.1. Defining the PAT-tuple

CMSSW version 2_2_9 was chosen as the base release for the samples, and the latest set of CVS tags and configuration for the PAT v1 was chosen as the base for the physics algorithms. This was the best and recommended PAT setup at the time this effort was started. In consultation with experts, some individual algorithms not present in the default PAT sequence were included and some existing algorithms updated to the latest versions. These were an improved photon ID, Jet Plus Track (JPT) jets, and track corrected MET. A fix for a bug in the treatment of saturated ECAL crystals also was included (requiring correction of the ECAL RecHits collection). It is worth noting that all of these features or bugs were fixed in CMSSW during the next release cycle.

Table 3 lists the event content, with the full list of collection names, number of items per event for each collection, and number of kB per event (averaged over 5000 LM0 events). This content was defined in consultation with members of the SUSY group.

The largest collections are the Monte Carlo information; the CaloTowers (to allow a more elaborate cleaning across collections); tracks; and the full list of Particle Flow (PF) candidates (for reconstruction of the PF objects). The next biggest collections are the four collections of jets that were added to these files (iterative cone 5, Siscone 5, PF jets, JPT).

The total event size could easily be reduced by selecting a subset of the Monte Carlo information to be kept, keeping only a subset of all tracks, not keeping the CaloTowers, etc. However, in this first round of production, the algorithms, samples and file content were chosen to be as inclusive as possible.

Table 3. Event content and size in the SUSY PAT-tuples (averaged over 5000 LM0 events). Collections are ordered by decreasing size. The listed content describes the current status of the PAT-tuples configuration and further refinements of it are planned for the next production iteration in CMSSW 3.1.

Collection name	Items/ev	kB/ev	%
recoGenParticles_genParticles__HLT	801.46	16.94	18.6
CaloTowersSorted_towerMaker__PAT	401.34	13.08	14.4
recoTracks_generalTracks__RECO	110.25	10.91	12.0
recoPFCandidates_particleFlow__RECO	334.04	10.70	11.8
patJets_allLayer1JetsSC5__PAT	19.78	10.36	11.4
patJets_allLayer1JetsIC5PF__PAT	21.24	4.75	5.2
patElectrons_allLayer1Electrons__PAT	1.62	3.92	4.3
patJets_allLayer1JetsIC5__PAT	15.46	3.85	4.2
patJets_allLayer1JetsIC5JPT__PAT	15.46	3.29	3.6
patPhotons_allLayer1Photons__PAT	4.54	3.24	3.6
patMuons_allLayer1Muons__PAT	1.49	1.84	2.0
recoSuperClusters_correctedHybridSuperClusters__RECO	10.53	1.53	1.7
recoGenJets_iterativeCone5GenJets__HLT	13.56	1.09	1.2
recoConversions_conversions__RECO	19.81	1.04	1.1
recoGenJets_sisCone5GenJets__HLT	12.93	1.02	1.1
recoSuperClusters_correctedMulti5x5SuperClustersWithPreshower__RECO	6.07	0.74	0.8
recoVertexs_offlinePrimaryVertices__RECO	1.06	0.73	0.8
recoTracks_ckfOutInTracksFromConversions__RECO	6.88	0.67	0.7
patMETs_allLayer1METsSC5__PAT	1.00	0.26	0.3
patMETs_allLayer1METsIC5__PAT	1.00	0.26	0.3
recoTracks_ckfInOutTracksFromConversions__RECO	1.95	0.19	0.2
recoPhotonIDs_PhotonIDProd_PhotonIDCutBasedProducer_PAT	4.54	0.14	0.2
patMETs_allLayer1METstcMET__PAT	1.00	0.10	0.1
patMETs_allLayer1METsPF__PAT	1.00	0.09	0.1
patTaus_allLayer1Taus__PAT	0.40	0.09	0.1
recoGenMETs_genMetNoNuBSM__HLT	1.00	0.04	0.0
recoGenMETs_genMet__HLT	1.00	0.04	0.0
recoPdfInfo_genEventPdfInfo__HLT	1.00	0.02	0.0
recoPhotonsToOnerecoPhotonIDsAssociation_PhotonIDProd_(…)_PAT	4.54	0.01	0.0
recoBeamSpot_offlineBeamSpot__RECO	1.00	0.01	0.0
int_genEventProcID__HLT	1.00	0.00	0.0
triggerTriggerEvent_hltTriggerSummaryAOD__HLT	1.00	0.00	0.0
Double_genEventWeight__HLT	1.00	0.00	0.0
Double_genEventScale__HLT	1.00	0.00	0.0
edmTriggerResults_TriggerResults__HLT	1.00	0.00	0.0
EventMetaData + EventHistory	1.00	0.29	0.3

The definition described above was the outcome of an iterative process with a small and dedicated group of analysers. It took around ten iterations and two weeks to settle on a final definition. During each iteration, small test files were produced, stored privately, and examined by the team. The final definition was documented on a single self-contained wiki page and served as a reference for PAT configuration in the SUSY group. It also served as a “recipe” that could be used to analyse the produced files, or for further private production of files.

5.2. Production and replication

The Monte Carlo samples to process were defined in wide consultation with members of the SUSY group, and were based on samples already used in previous analyses. Once the event samples were defined, a team of six people began to produce the requested samples using CRAB. A detailed recipe to produce and store the files at SUSY T2 sites was written. This included registering the files in a local instance of DBS. The production progressed very quickly at first, with more than 50 percent of the files being produced at the first attempt. However, a long tail in the production time was observed, which resulted in the whole production process taking between three and four weeks. The problems experienced chiefly were site-related, typically errors in reading from files or writing to storage elements at T2 sites. Generally, the response to problems was prompt. The success rate possibly could be improved by staging out the files at the same T2 where the production is run, if the data is located at the SUSY T2s. It should be noted, however, that users reported the exact same type of problems when accessing the produced files.

Table 4 shows details of the total size, number of events, and size per event for the 10 heaviest and 10 lightest samples (in terms of size per event). Sizes range from 45 kB to 120 kB (including Monte Carlo information), with an average of 80 kB/event, resulting in a total size of around 7 TB for all samples. In total, more than 100 M events were processed and stored at the SUSY T2 sites.

Table 4. Number of events, total size, and size per event for the 10 heaviest and 10 lightest PAT-tuple samples. The ttbar sample is also listed as a reference. Samples are ordered by decreasing size per event.

Dataset name	Number of events	Total size [GB]	Size per event [kB]
SUSY_LM8-sftsht	211302	23	118
BBJets1000toInf-madgraph	357618	37	111
Exotica_GMSB_GM1c	91171	9	110
Exotica_GMSB_GM1b	96760	10	109
QCD1000toInf-madgraph	1066863	110	108
SUSY_LM3-sftsht	153000	15	107
QCDpt3000	567040	58	107
Exotica_GMSB_GM1d	101193	10	107
QCDpt1400	584256	59	106
QCDpt800	2922476	290	104
TTbarJets-madgraph	946644	82	91
Zgg-madgraph	110150	5	52
Zgamma	106600	5	51
ZJets-madgraph	1262816	60	50
Wenu	1112967	53	50
ZinvisibleJets-madgraph	1018866	47	49
Wgamma	103122	4	49
QCDpt15	7938560	377	49
PhotonJetPt15	1035360	47	48
Wtaunu	1098500	49	47
WJets-madgraph	9745661	427	45

The final step in the production was to replicate the files produced to the five SUSY T2 sites, not only to ensure redundancy, but also to allow better user access. Since the files could not yet be registered in global DBS, the PhEDEx tool could not be used to perform the transfers between T2 sites. An interim solution was prepared by our colleagues from the Aachen group, to facilitate this until the registration in global DBS is possible.

Using this interim solution proved time consuming. The transfers between several of the sites were much slower than anticipated and only 75% of the requested transfer was completed after 4 weeks. Furthermore, the necessary bookkeeping had to be carried out by hand and thus was inefficient and error-prone. The SUSY signal samples were replicated to all five T2s and several of the background samples have been transferred successfully, too. However, some of the largest samples would take an unfeasibly long time with the interim solution and will remain at only one T2 site until the official tool is available. In general, the T2-T2 transfers (or T2-T1-T2) clearly need to be improved significantly.

5.3. Validation and evaluation

In order to collect feedback and validate the files produced for physics study, a team of volunteers was asked to test and record its findings on a set of blog-style wiki pages. Entries began during the event sample definition phase and continued after production. Known issues were recorded on the same self-contained wiki page as the definition of the event sample. Specifically, for this first attempt: some of the information was omitted from the files produced (L1 trigger information, PAT hemispheres, and MET muon correction vector). Currently, there are no other known issues, and none of these were considered as showstoppers.

Evaluating the entire procedure, lessons that may be learnt are that the validation of these samples should be integrated more tightly with the existing release validation of the SUSY group. The bottleneck in the production procedure was the sample replication between SUSY T2 sites. The tool to enter samples into global DBS and commissioning of T2 to T2 links will be of great importance in distributing group skims to the SUSY T2 sites. Generally, feedback from users has been positive, with only the limited availability of some of the larger samples being a common complaint.

In this case, the choice was made to produce PAT-tuples from RECO data samples. It would not have been possible to produce the PAT-tuples from AOD data samples, due to the lack of ECAL RecHits that were necessary for the saturation bug fix. The PAT data format provides more handles than the AOD format to reduce the event size, for example embedding relevant information in the output objects, although this necessarily comes with some loss of information. As a result, the PAT-tuple format is only slightly smaller than the AOD format (see Table 5). However, the PAT provides a data format that is more user-friendly and analysis-oriented than that of RECO objects, of which the AOD consists. Indeed, the PAT-tuples format is not just a subset of RECO collection but is a different data format (inheriting from RECO objects). This should be kept in mind when comparing the merits and limitations of AOD and PAT-tuples.

More specifically, the difference in sizes between AOD and PAT-tuple primarily comes from: the embedding of tau tag information, electron isolation deposits, and ECAL super-clusters; the drop of calorimeter taus (in favour of the Particle Flow taus); the reduction of number of jet collections that are saved (6 for the AOD, 4 for the PAT-tuple, including particle flow jets).

Table 5. Comparison between RECO, AOD and PAT-tuple event sizes for four different samples.

Dataset name	RECO event size [kB]	AOD[3] event size [kB]	PAT-tuple event size [kB]
SUSY_LM8-sftsht	670	180	120
SUSY_LM0-sftsht	560	150	100
TTbarJets-madgraph	510	136	91
WJets-madgraph	265	75	45

6. Summary of the SUSY group strategy

The general strategy for the data and workflow for SUSY analyses during the first physics run can be summarised in three categories:

1. The primary source for group skims will be Secondary Datasets. The SDs based on triggers are listed in Table 1 and should be made available to the SUSY group at T2 sites. Assuming that each of the important SUSY triggers will form a separate SD in RECO format, for the first 100 pb^{-1} of data, the size of these SDs will correspond to ~300 TB. Having these data easily accessible will speed up the production of group skims. In addition, it will help quick investigation of detector problems if need arises. In the case that there is no central production of the AOD, or if the AOD does not contain the information needed for all important SUSY analyses, the SUSY group would like to possess the flexibility to produce group skims using its own 100 kB analysis object that is tailored explicitly to the needs of the early searches.

2. As long as it is needed group skims will be derived from SDs with the data size reduction achieved by keeping only a 50 to 100 kB analysis object.
3. Once event size reduction is no longer sufficient to meet the resource constraints, additional cuts will be needed to reduce the size of the group skims to a sustainable level. The exact definition of these cuts will change with experience.

The structure and the content of the group skims are very inclusive and thus can be used for other physics analysis groups, which not only would result in some savings in disk space but also greatly enhance the coherence and consistency of the work in the different groups. This, however, requires that the physics groups agree on a common format for the 100 kB analysis object.

The SUSY group has had good experience with producing PAT-tuples. This format could be used for the SUSY group skims, but other formats, such as the AOD, could also be suitable.

The SUSY group is looking forward to a cross-physics group discussion on the format of the 100 kB analysis object and the development of a common data and workflow plan for physics analysis in CMS.

In the near future, it will be important to address the following questions:

- What is our current best specification of the content of the 100 kB analysis object?

3. This actually is the AODSIM content, which contains, in addition to the AOD, generator, generator jets for two algorithms, three additional jet collections, and some other collections of negligible size.

- What format will we use for the 100 kB analysis object (*e.g.*, AOD)?
- How will we perform the group skim replication: does the T2 to T2 transfer process scale? How well does the T2-T1-T2 replication procedures work and scale?
- What methods and samples are being developed with other CMS groups? Is there a global approach that meets the needs of many groups and that would reduce the work load on individual groups?

7. Acknowledgements

We thank M. Giffels and T. Kress for providing performance details of the T2 to T2 data replication procedure.

8. References

[1] https://twiki.cern.ch/twiki/bin/view/CMS/SusyPagReferenceAnalyses.

[2] https://twiki.cern.ch/twiki/bin/view/CMS/TriggerMenuDescription1E31Devel Description of the trigger table for startup. It should be noted that the start-up trigger table menus are still under development and therefore some of the thresholds and trigger definitions presented here may still be further refined.

[3] https://twiki.cern.ch/twiki/bin/view/CMS/SusyPatLayer1