



EOSDIS

NASA'S EARTH OBSERVING SYSTEM
DATA AND INFORMATION SYSTEM

Cloud onboarding with NGAP

Cloud Onboarding Session
ESIP Summer Meeting 2017
11am July 27th 2016

Doug Newman (NASA Earthdata – Raytheon)
Dan Pilone (NASA Earthdata – Element 84)

The material is based upon work supported by the National Aeronautics and
Space Administration under Contract Number **NNG15HZ39C**

DN-SESIP-0717

Agenda

- What is NGAP?
- Two important things
- How things do (and don't) change
 - Deployment
 - Configuration
 - Security
 - External communications
 - Scaling
 - Logging
 - Monitoring
 - Metrics
 - Contingency and Recovery
- How did it go with the Common Metadata Repository?

Dan Pilone

WHAT IS NGAP?

- Formulation
- Implementation
- Primary Ops
- Extended Ops



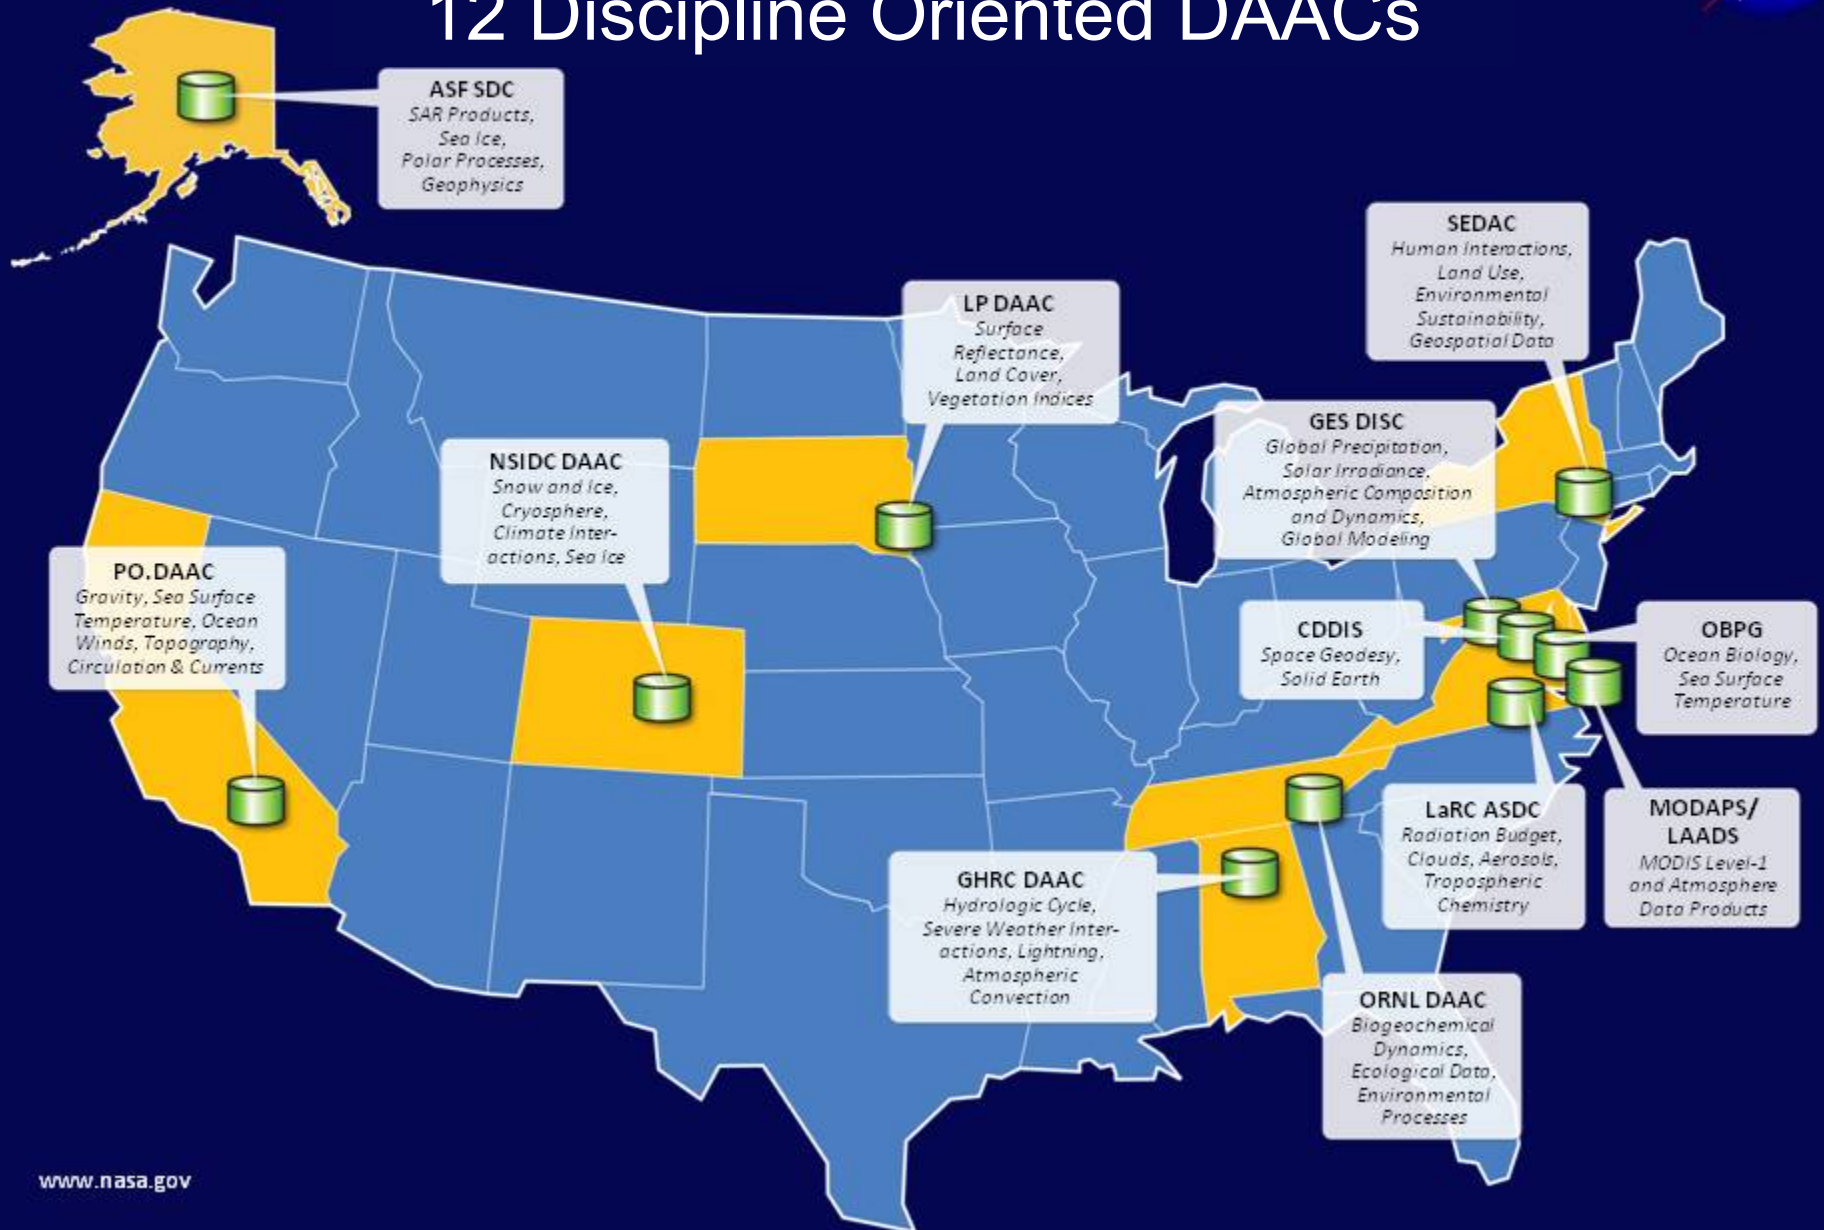
Earth Science Instruments on ISS:


RapidScat, CATS,
LIS, SAGE III (on ISS), TSIS-1, OCO-3, ECOSTRESS,
GEDI, CLARREO-PF






12 Discipline Oriented DAACs





Background: Current and Estimated Archive Growth



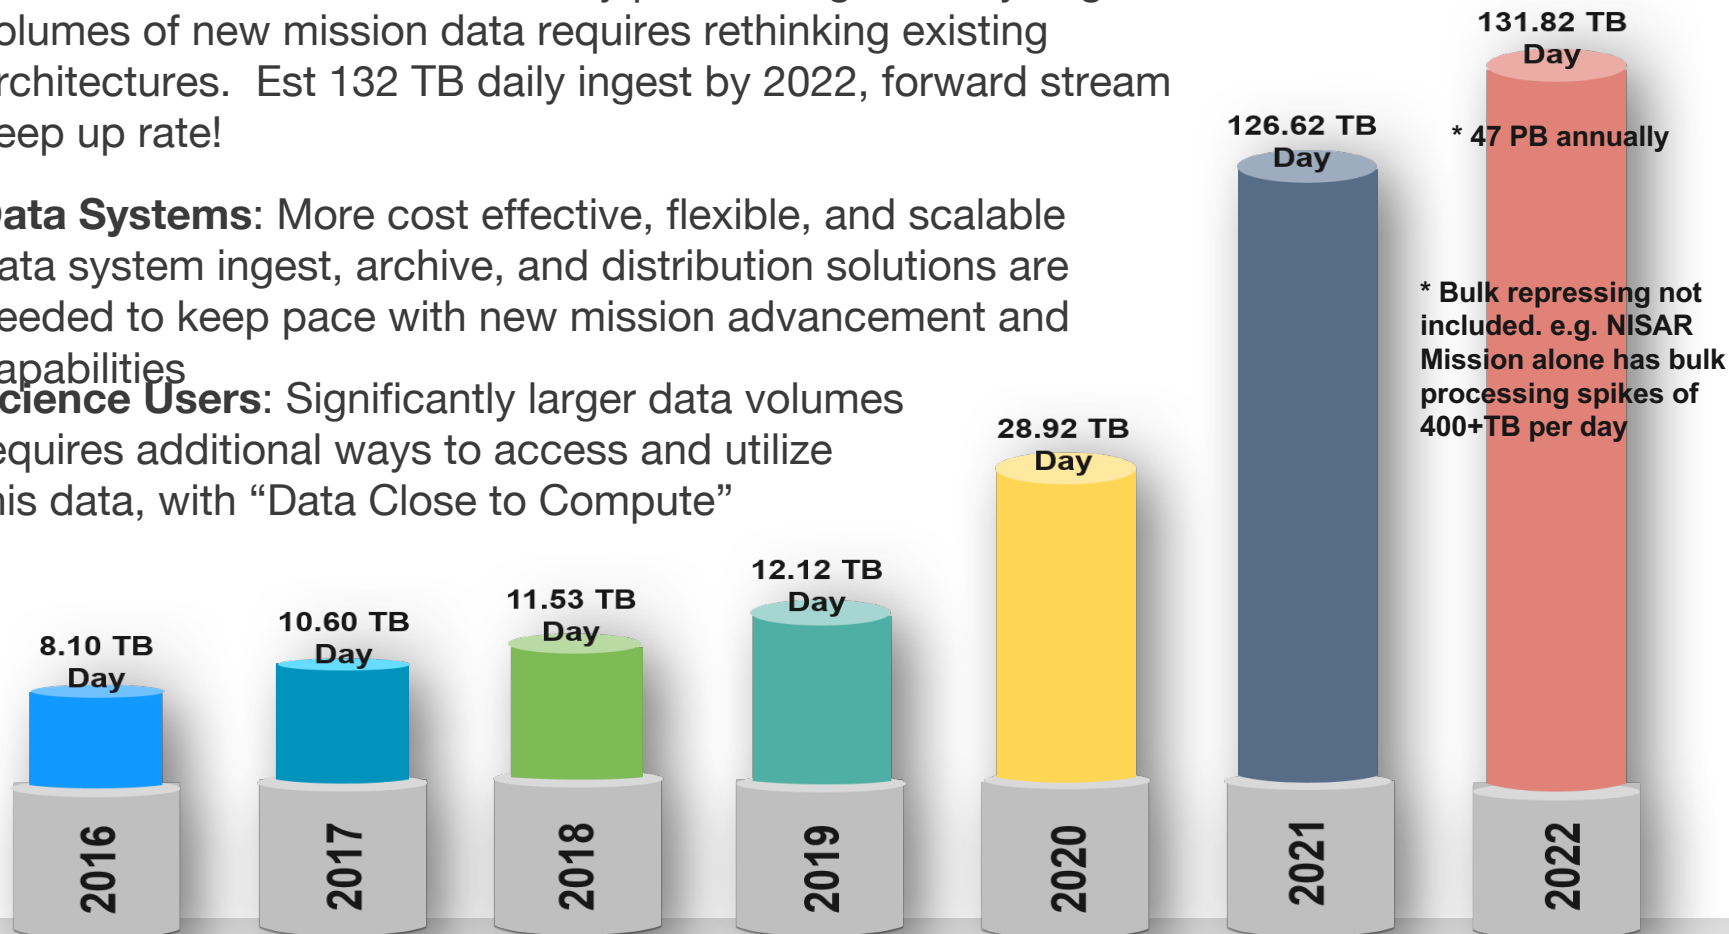


Background: Motivation for Cloud

Growth of Mission Data & Processing: Projected rapid archive growth and the need to effectively process significantly larger volumes of new mission data requires rethinking existing architectures. Est 132 TB daily ingest by 2022, forward stream keep up rate!

Data Systems: More cost effective, flexible, and scalable data system ingest, archive, and distribution solutions are needed to keep pace with new mission advancement and capabilities

Science Users: Significantly larger data volumes requires additional ways to access and utilize this data, with “Data Close to Compute”





ExCEL Efforts and Project Prototypes



NGAP

NASA Compliant General Application Platform (NGAP), an operational, dev-ops, and sandbox AWS cloud based operating environment.

ASF WOS Prototype

AWS/NGAP Web Object Storage (WOS) prototyping large volumes of mission data dynamically between AWS S3, S3-IA, and Glacier object storage. Managed out of Alaska Satellite Facility

Earthdata Search Client to Cloud

NASA Earth Science data search by keyword and advanced filters such as time and space

Cumulus

Prototype addressing core EOSDIS capabilities including data ingest, archive, management, and distribution of large volumes of EOS data.

Getting Ready for NISAR (GRFN)

Integrated prototype of science product generation and delivery from a DAAC system focused on coupling ASF DAAC and JPL ARIA systems.

CATEES

Easy-to-use Python tools packaged to support EOSDIS cross-DAAC science workflows and analytics over large volumes of EOS data in AWS.

ECC to Cloud Study

Earth Code Collaborative (ECC) study to determine cloud ready capabilities to migrate into AWS/NGAP platform.



ExCEL Efforts and Project Prototypes Continued

GIBS in the Cloud

Migrating GIBS to the AWS/NGAP Cloud based on recommendations made in the “GIBS in the Cloud Study”

Earthdata Login to Cloud Study

Study to determine and recommend migrating the Earthdata Login into AWS/NGAP cloud environment

CMR to Cloud

Migration of the Common Metadata Repository, into the AWS/NGAP platform based on recommendations made in the CMR to Cloud study.

OPeNDAP/HDF Cloud Studies

Study to determine and recommend a cloud native integration of OPeNDAP accessing HDF5 and netCDF4 data on AWS/NGAP platform.

NEXUS

Prototype to accelerate end-user analysis of remote sensing data, highly parallel to better enable science discovery

Network Prototypes

Network prototypes to support to test security, monitoring, logging, and to perform R&D testing to support all ExCEL project prototypes.





(01) Full Scale Deployment (?)

Full scale enterprise deployment of EOSDIS services and infrastructure to the cloud

01

02

(02) Partial Deployment (?)

Select deployment of EOSDIS services and/or infrastructure to the cloud

(03) Cloud Stand-down (?)

No EOSDIS services or infrastructure operationally migrated to the cloud

03

(04) Decision Point (?)

More prototyping required, or cloud hybrid, or other next steps based on ExCEL prototyping and business analysis results

04

Determining Project Success

Project success is determined by viable outcomes of fully completed project prototypes and business analysis.

- or -

Technical and business results of the ExCEL project needed for strategic decision on EOSDIS and the cloud.





ExCEL Technical Foundational Capabilities


Data Management:

Development/execution of information lifecycle needs of NASA mission based earth science data sets

Data Ingest: Multi-mission, multi-disciplines data ingest into the cloud

Data Distribution:

Distribution includes raw data, sub-setting, visualization, and more



Data Archive: Preserve and protect NASA mission based earth science data assets

Processing: Processing for Science Data Systems (SDS) and Cloud based on-the-fly data analysis and API access to analytic optimized cloud storage

Platform: Cloud infrastructure capabilities on Amazon Web Services (AWS) and the EOSDIS NASA Compliant General Application Platform, including networking and IT Security




EOSDIS

Cloud Evolution (ExCEL) Project

DN-SESIP-0717



ExCEL Cloud Platform Capabilities



ExCEL Cloud Platform Capabilities

Cloud Platform Service Capabilities

All components of a centralized and compliant cloud infrastructure. This includes robust/flexible networking capabilities, tactical/compliance IT Security capabilities, and cloud native computing and storage distribution capabilities

Partnerships and Access to Commercial Cloud

Joint Cloud Service Program Office (CSPO) provides IT Security support, cloud infrastructure support, business level support, and contractual support to Amazon Web Services (AWS)

Test and Operational Cloud Platforms

NASA's Compliant General Application Platform (NGAP) is built off CSPO's General Purpose Managed Cloud Environment (GPMCE) and provides test/operational cloud based Platform-as-a-Service (PaaS) and Infrastructure-as-a-Service (IaaS) for EOSDIS

NASA Compliant General Application Platform (NGAP)

What is NGAP?

NGAP is the **NASA Compliant General Application Platform**. It provides a cloud-based Platform-as-a-Service (PaaS) and Infrastructure-as-a-Service (IaaS) for ESDIS applications.

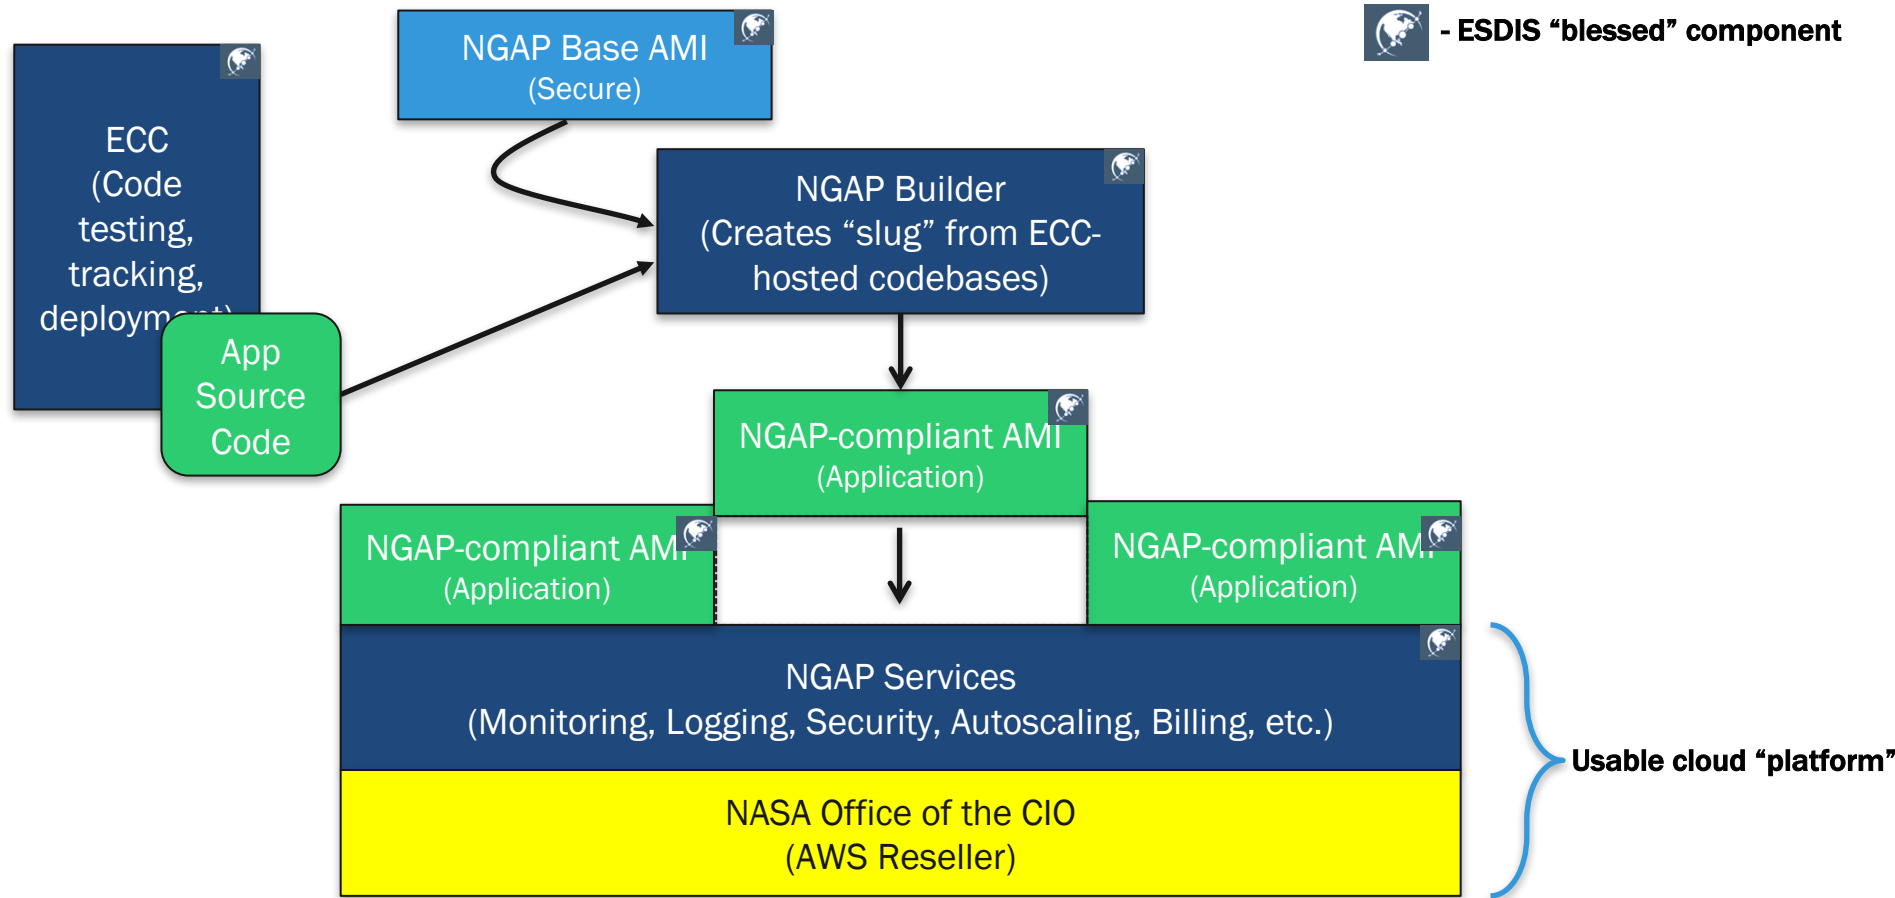
NGAP as a Platform

NGAP Services

(Monitoring, Logging, Security, Autoscaling, Billing, etc.)

NASA's Office of the Chief Information Officer
(AWS Reseller)

Layer compliance throughout the architecture

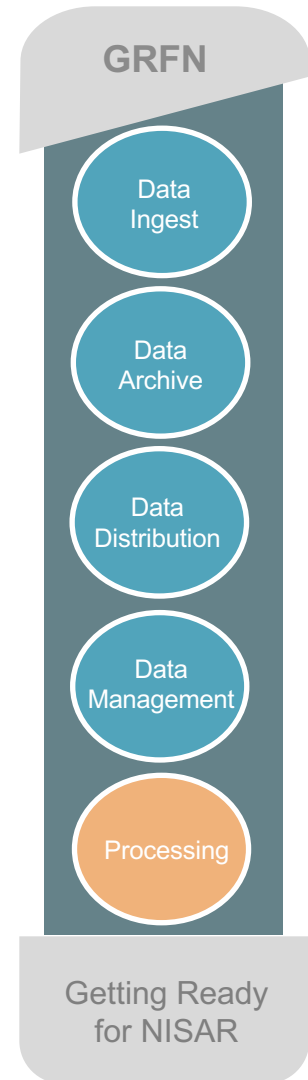


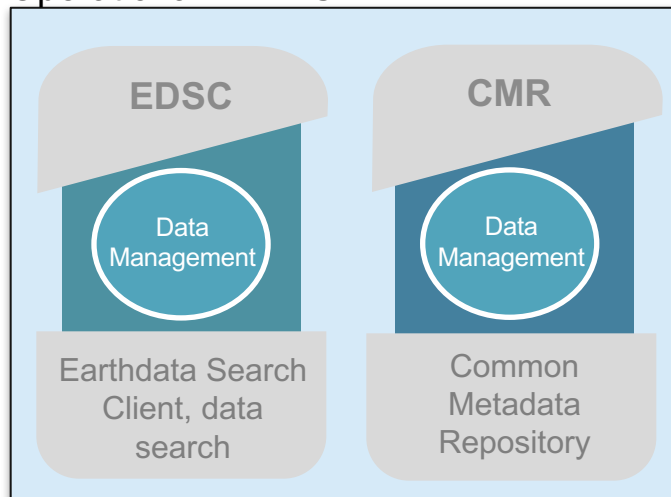




Background Technical Core Capability Mapping to ExCEL Prototypes

Mapping EOSDIS technical foundational capabilities to ExCEL project prototype activates

Operational in AWS



DN-SESIP-0717

Doug Newman

ONBOARDING EXPERIENCES WITH NGAP

TWO IMPORTANT THINGS

1. Instances are ephemeral

NGAP deployments follow a blue-green deployment process

To maximize the availability and performance of our applications, a deployment is spun up in parallel with the existing deployment. When the secondary deployment is ready it swaps with the existing deployment which is then discarded

NGAP application instances are not available in perpetuity

2. No ssh

To preserve the integrity of an application instance, ssh (secure shell) access is limited to NGAP personnel

HOW THINGS DO (AND DON'T) CHANGE

Deployment

Bamboo is used to perform deployments

SIT deployments are controlled by the **development** team

Production and UAT deployments are tightly controlled by the **DEVOPS** team

Earthdata Operations maintain a [Deployment Doctrine](#) that is publicly available

Configuration

12-factor-app practices encourage the storage of configuration with the environment rather than the code

We developed the Earthdata Environment Configuration Service (**EECS**) to configure our applications

EECS provides an API to read and write JSON-formatted configuration for your application on a per-environment basis

If an implementer chooses not to use EECS then configuration should be externalized from code

Security (1 of 2)

The responsibility for identifying and resolving security issues and software patches rests with the GP-MCE

They will release Amazon Machine Instances (AMIs) to NGAP

NGAP will release that AMI to NGAP PROD after SIT and UAT testing

The application team will deploy the new AMI with any deployment of their applications that exist in NGAP PROD

This approach has a number of elements that need to be allowed for

Security (2 of 2)

1. Not all applications have a presence in NGAP SIT and UAT
2. Once an AMI hits NGAP PROD all deployments there will use the new AMI*

*We plan to mitigate this by giving an operator choices in AMI at certain points

External communications

On-premises solutions generally have a static set of IP addresses that an external entity can expect traffic from

NGAP instances are ephemeral

NGAP applications have a range of possible IP addresses

Stick to standard ports if possible. Amazon Web Services (AWS)/GP-MCE/NGAP do not block outgoing traffic to standard ports.

Scaling

Manual scaling is extremely simple to achieve via the ngap-cli application.

```
> bundle exec ngap ps:scale <app name> 2
```

Logging (1 of 2)

No ssh access. No log files.

NGAP automatically generates all needed artifacts to analyze application and access logs with **Splunk**

Use Splunk

Logging (2 of 2)

Search | Splunk 6.5.3

https://logs.earthdata.nasa.gov/en-US/app/search/search?earliest=-15m&latest=now&q=search sourcetype%3Dprod_ngap* sourcetype%3Dprod_ngap_app_cm-search-prod&display.page.s

Search & Reporting

New Search

sourcetype="prod_ngap_app_cm-search-prod"

69,189 events (6/19/17 10:34:50.000 AM to 6/19/17 10:49:50.000 AM) No Event Sampling

Events (69,189) Patterns Statistics Visualization

Format Timeline Zoom Out Zoom to Selection Deselect

1 minute per column

List Format 20 Per Page

< Hide Fields All Fields

Selected Fields

- host 10
- source 1
- sourcetype 1

Interesting Fields

- index 1
- linecount 2
- splunk_server 1

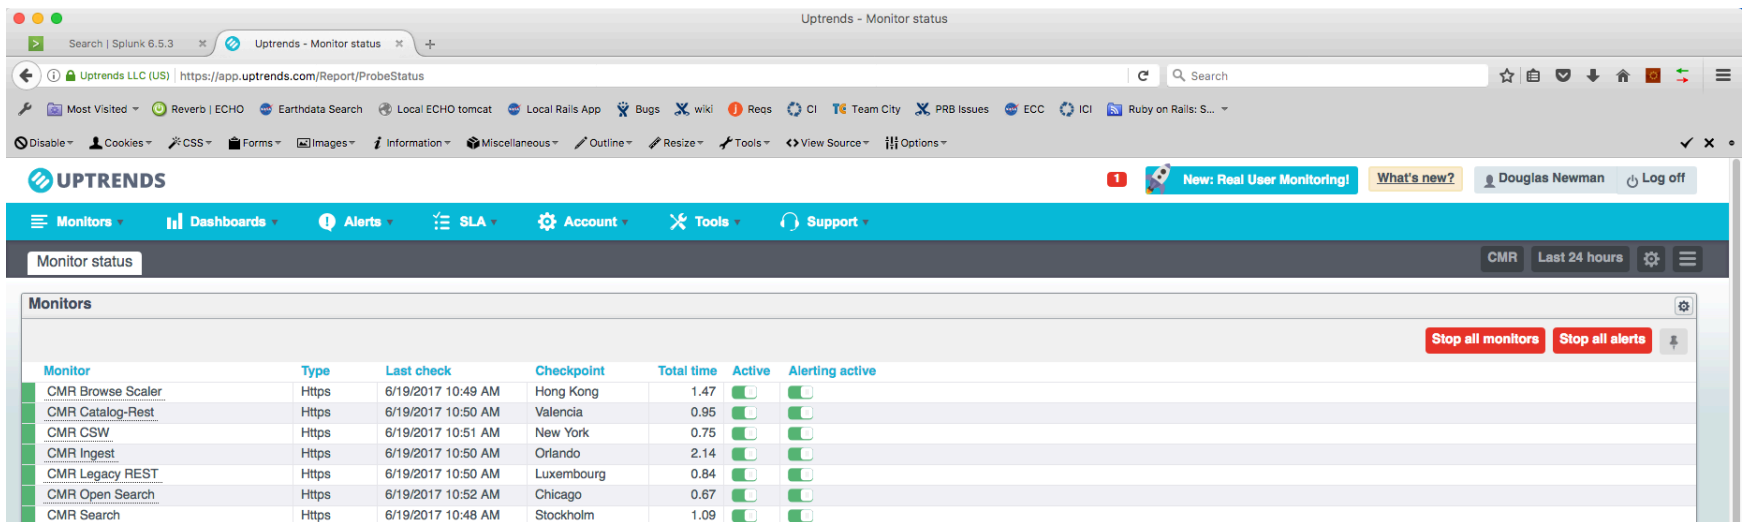
+ Extract New Fields

i	Time	Event
>	6/19/17 10:49:49.522 AM	web.ea00fa3 2017-06-19 14:49:49.522 ip-198-122-13-25.aws.nasa.gov [8b6e563e-575b-4463-8ba6-5d10d8e484b6] INFO [cmr.search.services.query-service] - Found 2 collections in 9 ms in format :echo10 with aql: <query><for value="collections" /><dataCenterId><value>LPDAAC_ECS</value></dataCenterId><where><collectionCondition><dataSetId><list><value>ASTER L1A Reconstructed Unprocessed Instrument Data V003</value><value>ASTER Level 1 precision terrain corrected registered at-sensor radiance V003</value></list></dataSetId></collectionCondition></where></query>.
>	6/19/17 10:49:49.522 AM	web.ea00fa3 2017-06-19 14:49:49.522 ip-198-122-13-25.aws.nasa.gov [8b6e563e-575b-4463-8ba6-5d10d8e484b6] INFO [cmr.common-app.services.search] - query-execution-time: 6 result-gen-time: 0 host = ngap.prod.app.cmr-search-prod-i-0eb4bc99a3e365df:ip-198-122-13-25 source = /ngap/log/app.cmr-search-prod.log sourcetype = prod_ngap_app_cm-search-prod
>	6/19/17 10:49:49.522 AM	web.ea00fa3 2017-06-19 14:49:49.522 ip-198-122-13-25.aws.nasa.gov [8b6e563e-575b-4463-8ba6-5d10d8e484b6] INFO [cmr.search.data.metadata-retrieval.metadata-cache] - get-formatted-concept-revisions of 2 concepts total: 0 get-cached-metadata-in-format 0 revision-format-maps->concepts: 0 transform-and-cache: 0 fetch-and-cache: 0 fetch: 0 order-concepts: 0 host = ngap.prod.app.cmr-search-prod-i-0eb4bc99a3e365df:ip-198-122-13-25 source = /ngap/log/app.cmr-search-prod.log sourcetype = prod_ngap_app_cm-search-prod
>	6/19/17 10:49:49.522 AM	web.ea00fa3 2017-06-19 14:49:49.522 ip-198-122-13-25.aws.nasa.gov [8b6e563e-575b-4463-8ba6-5d10d8e484b6] INFO [cmr.common-app.services.search.elastic-search-index] - Elastic query took 1 ms. Connection elapsed: 6 ms host = ngap.prod.app.cmr-search-prod-i-0eb4bc99a3e365df:ip-198-122-13-25 source = /ngap/log/app.cmr-search-prod.log sourcetype = prod_ngap_app_cm-search-prod
>	6/19/17 10:49:49.516 AM	web.ea00fa3 2017-06-19 14:49:49.516 ip-198-122-13-25.aws.nasa.gov [8b6e563e-575b-4463-8ba6-5d10d8e484b6] INFO [cmr.common-app.services.search.elastic-search-index] - Executing against indexes [collection_search_alias] the elastic query: {query:{filtered:{query:{match_all:{}}},:filter:{and:{filters:{term:{provider-id "LPDAAC_ECS"}}}:terms:{entry-title ("ASTER Level 1 precision terrain corrected registered at-sensor radiance V003" "ASTER L1A Reconstructed Unprocessed Instrument Data V003")}:execution "plain"}}}} with sort ({:_script {:_script "/* Elastic groovy script to calculate the temporal overlap of a collection over the input ranges.\n Collection start and end dates are doc['start-date'] and doc['end-date'].\n For each temporal range, calculate the amount of overlap for the collection. Add them up\n and divide by the sum of the span of the ranges (rangeSpan) to get the overall overlap.\n*/\n\n def totalOverlap = 0;\n\n for (range in temporalRanges)\n {\n def overlapStartDate = range.start_date;\n if (doc['start-date'].empty == false && doc['start-date'].value > overlapStartDate)\n {\n overlapStartDate = doc['start-date'].value;\n }\n\n def overlapEndDate = range.end_date;\n if (doc['end-date'].empty == false && doc['end-date'].value < overlapEndDate)\n {\n overlapEndDate = doc['end-date'].value;\n }\n\n if (overlapEndDate > overlapStartDate)\n {\n totalOverlap += overlapEndDate - overlapStartDate;\n }\n\n if (totalOverlap / rangeSpan > 0) {totalOverlap / rangeSpan};\n\n else{\n 0; /* Temporal overlap is 0 */\n }\n\n :type :number, :params {temporalRanges {}, :rangeSpan 0, :order :desc};\n :usage-relevancy-score {order :desc, :missing 0};\n :short-name {order "asc"}; :version-id {order "desc"}; :concept-seq-id {order "asc"}; :revision-id {order "desc"}}} with aggregations nil and highlights nil host = ngap.prod.app.cmr-search-prod-i-0eb4bc99a3e365df:ip-198-122-13-25 source = /ngap/log/app.cmr-search-prod.log sourcetype = prod_ngap_app_cm-search-prod
>	6/19/17 10:49:49.512 AM	web.ea00fa3 2017-06-19 14:49:49.513 ip-198-122-13-25.aws.nasa.gov [8b6e563e-575b-4463-8ba6-5d10d8e484b6] INFO [cmr.search.api.routes] - Searching for concepts from client CMR-Legacy-Services in format :echo10 with aql: <query><for value="collections" /><dataCenterId><value>LPDAAC_ECS</value></dataCenterId><where><collectionCondition><dataSetId><list><value>ASTER L1A Reconstructed Unprocessed Instrument Data V003</value><value>ASTER Level 1 precision terrain corrected registered at-sensor radiance V003</value></list></dataSetId></collectionCondition></where></query>.

External monitoring (1 of 2)

External monitoring strategies **will not** be affected by the transition to NGAP

Monitoring of public APIs and applications do not change



The screenshot shows the Uptrends web interface. The top navigation bar includes links for Monitors, Dashboards, Alerts, SLA, Account, Tools, and Support. The main content area is titled 'Monitor status' and displays a table of active monitors. The table has columns for Monitor, Type, Last check, Checkpoint, Total time, Active, and Alerting active. There are two red buttons at the top right of the table: 'Stop all monitors' and 'Stop all alerts'.

Monitor	Type	Last check	Checkpoint	Total time	Active	Alerting active
CMR Browse Scaler	Https	6/19/2017 10:49 AM	Hong Kong	1.47	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CMR Catalog-Rest	Https	6/19/2017 10:50 AM	Valencia	0.95	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CMR CSW	Https	6/19/2017 10:51 AM	New York	0.75	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CMR Ingest	Https	6/19/2017 10:50 AM	Orlando	2.14	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CMR Legacy REST	Https	6/19/2017 10:50 AM	Luxembourg	0.84	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CMR Open Search	Https	6/19/2017 10:52 AM	Chicago	0.67	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CMR Search	Https	6/19/2017 10:48 AM	Stockholm	1.09	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Internal monitoring (1 of 3)

Internal monitoring strategies **may** be affected by the transition to NGAP

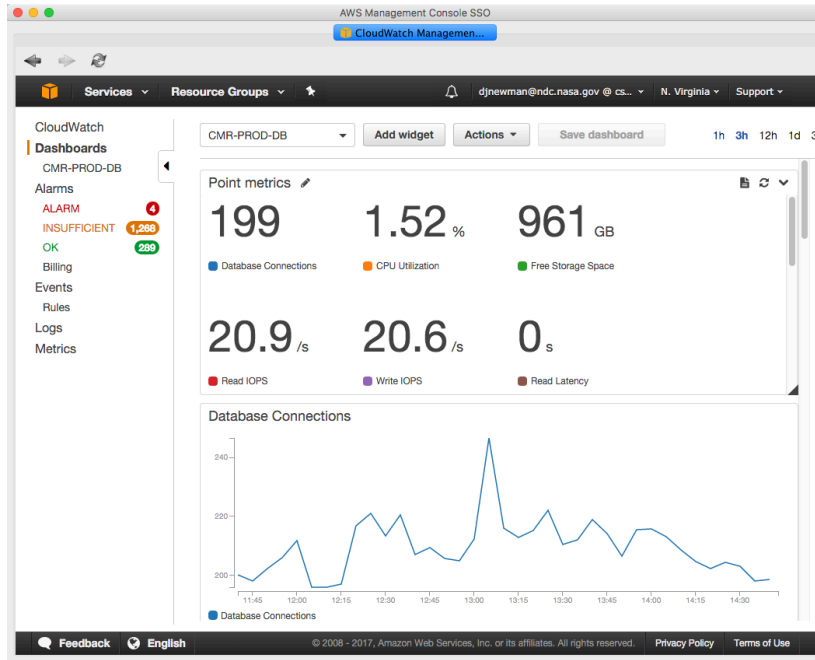
No ssh access.

Monitoring must be done using one of the following methods,

Internal monitoring (2 of 3)

- NGAP generates metrics, alarms and notifications for our ephemeral instances
 - Disk utilization
 - CPU utilization
 - Memory utilization
- Custom alarms may also add alarms to metrics associated with static AWS resources

Internal monitoring (3 of 3)



The screenshot displays the AWS Management Console for the 'EC2 Management Console'. The left sidebar shows navigation options: EC2 Dashboard, Events, Tags, Reports, Limits, INSTANCES, Spot Requests, Reserved Instances, Scheduled Instances, Dedicated Hosts, IMAGES, AMIs, Bundle Tasks, ELASTIC BLOCK STORE, Volumes, Snapshots, NETWORK & SECURITY, Security Groups, and Elastic IPs. The main content area shows a table of EC2 instances.

Name	Instance ID	Instance Type	Availability Zone	Instance State
NGAP prod App cmr-search-prod 4462	i-06e771b62d3b4e78f	m4.large	us-east-1a	running
NGAP prod App cmr-search-prod 4462	i-0949a0971e1ad2c6b	m4.large	us-east-1b	running
NGAP prod App cmr-search-prod 4462	i-0ae33c968d99b4c8f	m4.large	us-east-1b	running
NGAP prod App cmr-search-prod 4462	i-0bdb7eb3e9611901f	m4.large	us-east-1b	running
NGAP prod App cmr-search-prod 4462	i-0d5c2d44f96a40a94	m4.large	us-east-1a	running

Metrics

Metrics can be obtained using the following applications aligned with NGAP,

- Splunk
- AWS CloudWatch*

And external applications such as,

- Google Analytics
- Uptrends

These can be leveraged during issue triage, reporting and performance analysis

*We are looking into piping AWS metrics into Splunk

DN-SESIP-0717

Contingency & Recovery

- NGAP can currently deploy our applications to multiple availability zones (AZ) within the US-East region for Platform as a service (PaaS) applications
- In the future, we could support deployment across multiple regions (within CONUS)
- If one AZ goes down the other one is still there. Our applications keeps working
- We expect to be able to leverage recovery capabilities provided by the cloud and NGAP


HOW DID IT GO WITH CMR?

CMR?

‘The Common Metadata Repository (CMR) is a high-performance, high-quality, continuously evolving metadata system that catalogs Earth Science data and associated service metadata records’


- 33K collections
- 380 million granules
- 95% of queries are resolved in less than 1 second
- 12 node elastic search cluster (1.4 TB) for search
- Oracle Relation Database Service (RDS) for metadata persistence
- 14 micro services
 - On premises – 5 hosts (1 instance on each)
 - NGAP – 42 application instances (varying numbers of redundancy)

Performance (1 of 2)



DN-SESIP-0717

Performance (2 of 2)



DN-SESIP-0717

Stability

- 2017 Prod uptime on-premises: 99.83%
- 2017 Prod uptime on-cloud: 99.67%
- 2017 UAT uptime on-premises: 99.77%
- 2017 UAT uptime on-cloud: 99.77%
- 2017 SIT uptime on-premises: 98.32%
- 2017 SIT uptime on-cloud: 99.65%

This material is based upon work
supported by the National
Aeronautics and Space
Administration under Contract
Number **NNG15HZ39C.**

Raytheon