ata-driven pipelines high performance computing machine and deep learning relationship an
raph analytics streaming data analytics iot technology consulting high performance analytic
rchitecture learned and intelligent systems technology vision, strategy, and practice blockchai
nd smart contracts data ontology and design data-driven pipelines high performance computin
achine and deep learning relationship and graph analytics streaming data analytics io
echnology consulting high performance analytics architecture infrastructure intelligent system
eal-time event-driven systems blockchain and smart contracts data ontology and design data
riven pipelines high performance computing machine and deep learning relationship and grap
nalytics streaming data analytics iot technology consulting high performance analytic
rchitecture infrastructure intelligent systems event-driven real-time analytics blockchain an
mart contracts data ontology and design data-driven pipelines high performance computin
achine messaging topologies relationship and graph analytics streaming data analytics io
echnology consulting high performance analytics architecture infrastructure intelligent system
eal-time event-driven systems blockchain and smart contracts data ontology and design data
riven pipelines high performance computing machine and deep learning relationship and grap
nalytics streaming data analytics iot technology consulting high performance analytic
rchitecture infrastructure intelligent systems real-time event-driven systems blockchain an

Providentia Worldwide

(Un)Obscured By Clouds

Applying Cloud Techniques To Address Complexity In High Performance System Integrations

Argonne NATIONAL LABORATORY

HPC USER FORUM

Arno Kolster, Principal & Co-Founder

Providentia Worldwide

What Prompted This Talk?

Integrations between HPC and Enterprise have been coming for quite some time.

HPC in the cloud is now a reality - easier and cheaper to use than ever before.

FOMO on latest technical trends, technology itself and desire to increase ROI.

Majority of our work has been leveraging our hyperscale/cloud experience in HPC environments - time to talk about ’eating our own dog food’.

But cloud still a huge learning curve and paradigm shift for traditional HPC.

As it is for enterprise embracing HPC tech – i.e. GPUs for ML/DL/AI.

Use Case #1

Leadership Class Supercomputer Operations Integration

“As we build faster and faster supercomputers, their energy efficiency becomes more and more important,” said Jim Rogers, computing and facilities director for ORNL’s National Center for Computational Sciences. “We wanted to couple Summit’s mechanical cooling system with its computational workload to optimize efficiency, which can translate to significant cost savings for a system of this size.”

Drivers For Summit Cooling Intelligence

Summit is the largest water cooled supercomputer in the world.

Minor fluctuations in temperature across 4,608 nodes can have a dramatic positive or negative impact on Summit’s operational efficiency.

No mechanism exists to monitor the temperature in real time so that systems operators can make adjustments as appropriate.

Can take upwards of 3 mins for change in water temp to reach furthest node in the grid.

Systems-level intelligence can inform decisions and impact costs, efficiency and operational excellence.

SCI Design Goals

Architect, develop and implement a scalable platform to:

Ingest line-rate telemetry from Summit & supporting systems in real time.

Add data sets easily to enhance analytics and insight.

Employ cloud-style ‘reliable messaging’ as a common interface to provide consistency and coherency for ‘in-flight’ data.

Scale data at orders-of-magnitude without proportional staff increases.

Obey the axiom: *No single failure results in a single point-of-failure.*

Currently 4 Source-Data Metrics

1. IBM OpenBMC framework (99 metrics * 4608 nodes/second = 456,192/sec)
2. IBM LSF jobs data for running applications (~10sec)
3. NOAA weather/wet bulb for Oak Ridge area (continuous)
4. Programmable Logic Circuit (PLC) water flow (continuous)

Instrumentation and monitoring data for developed applications (ad hoc)

Summarized and collated metrics (continuous)

Instrumented Ascent (dev cluster) at the same time using same platform

TOTAL ~460,000 metrics per second available for real-time analytics

Providentia Worldwide

SCI Architecture Created

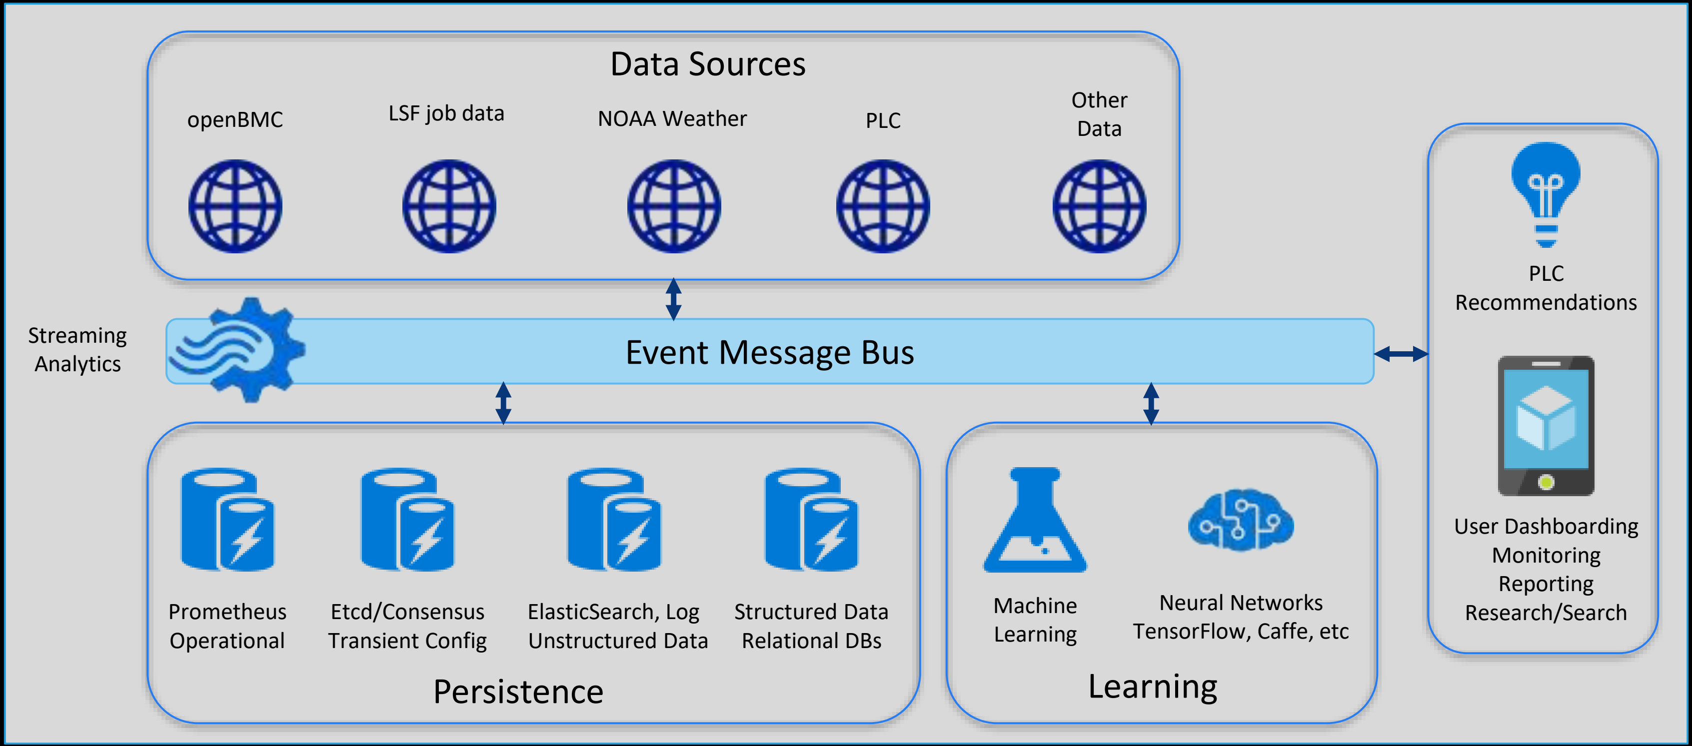

Providentia Worldwide

Technologies Used & Integrated

High-availability *microservices* designed for scalable, small footprint

Kafka message bus for ubiquitous message delivery

Docker containers for ease of deployment and elastic scalability

Openshift/Kubernetes for container orchestration and instrumentation

Spark Streaming for *on-the-wire* data analytics

Prometheus for service discovery metrics collection and time-series db

Jupyter Notebooks for models and data science exploration (Python/Spark)

Grafana, Seaborn for visualization

Breakthroughs

Created a new paradigm for analyzing data beyond data-at-rest or data at the sensor.

Data-in-flight is agnostic and can be captured prior to being translated into specific software formats, some of which can be proprietary.

Overlapping data sets from disassociated sources can now be displayed on an identical scale (time), which is something new.

What Can Be Done Now

Capture 99 node metrics per second for streaming analytics; for example,

- Core temps for CPU, GPU, memory DIMMs, HBMs
- Temps by CPU core, GPU core
- Power to the node, fans and to other individual components

Visualize per-second node temperature+power and temperature+jobs (currently).

Analyze behavior from both a systems level and node level.

Optimize operations by combining metrics from Summit, control systems & weather.

Add more real-time data streams from varying sources to augment analytics.

Impact

Real-time analytics helps operators detect anomalies quickly and in a manner that is more responsive than querying a data warehouse. This approach is especially valuable with such a large investment in a critical asset.

Real-time data can ultimately help micro-manage such a large system in ways that are far more efficient and cost effective than human-only intervention.

- Efficiencies in Summit ecosystem, job throughput and individual jobs
- Micro-managed cooling leads to running at higher temps with less risk
- Greater efficiency leads to Summit optimization, health and lower costs

This platform captures all the data needed to conduct predictive analytics, thus ensuring new levels of insight, alerting and monitoring.

Providentia Worldwide

Real-Time Visualization

Summit Power/Jobs Every Second

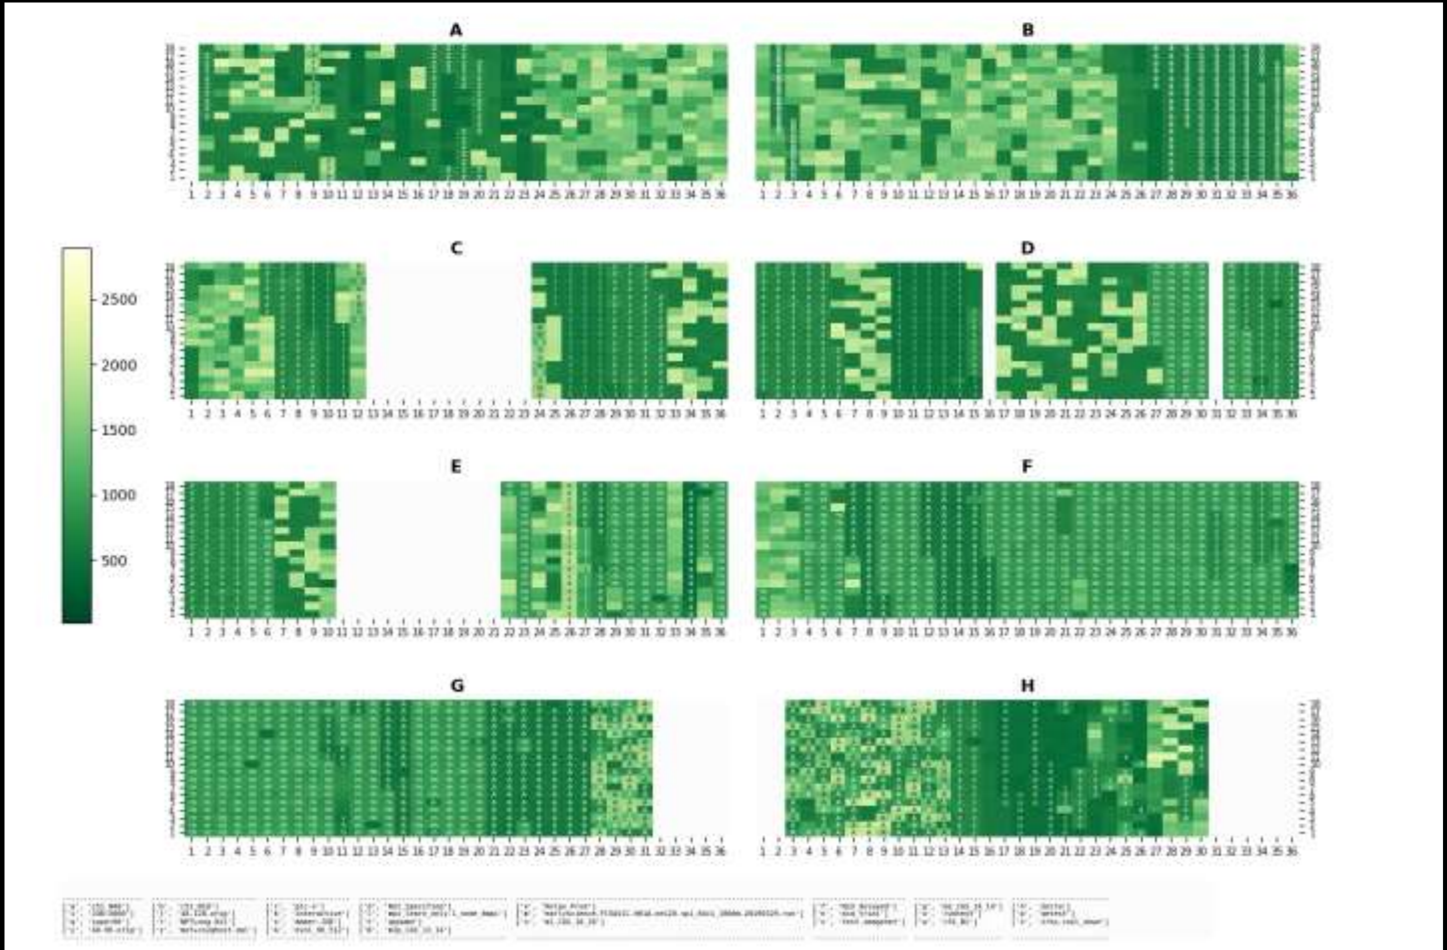

Summit Temperature/Jobs Every Second

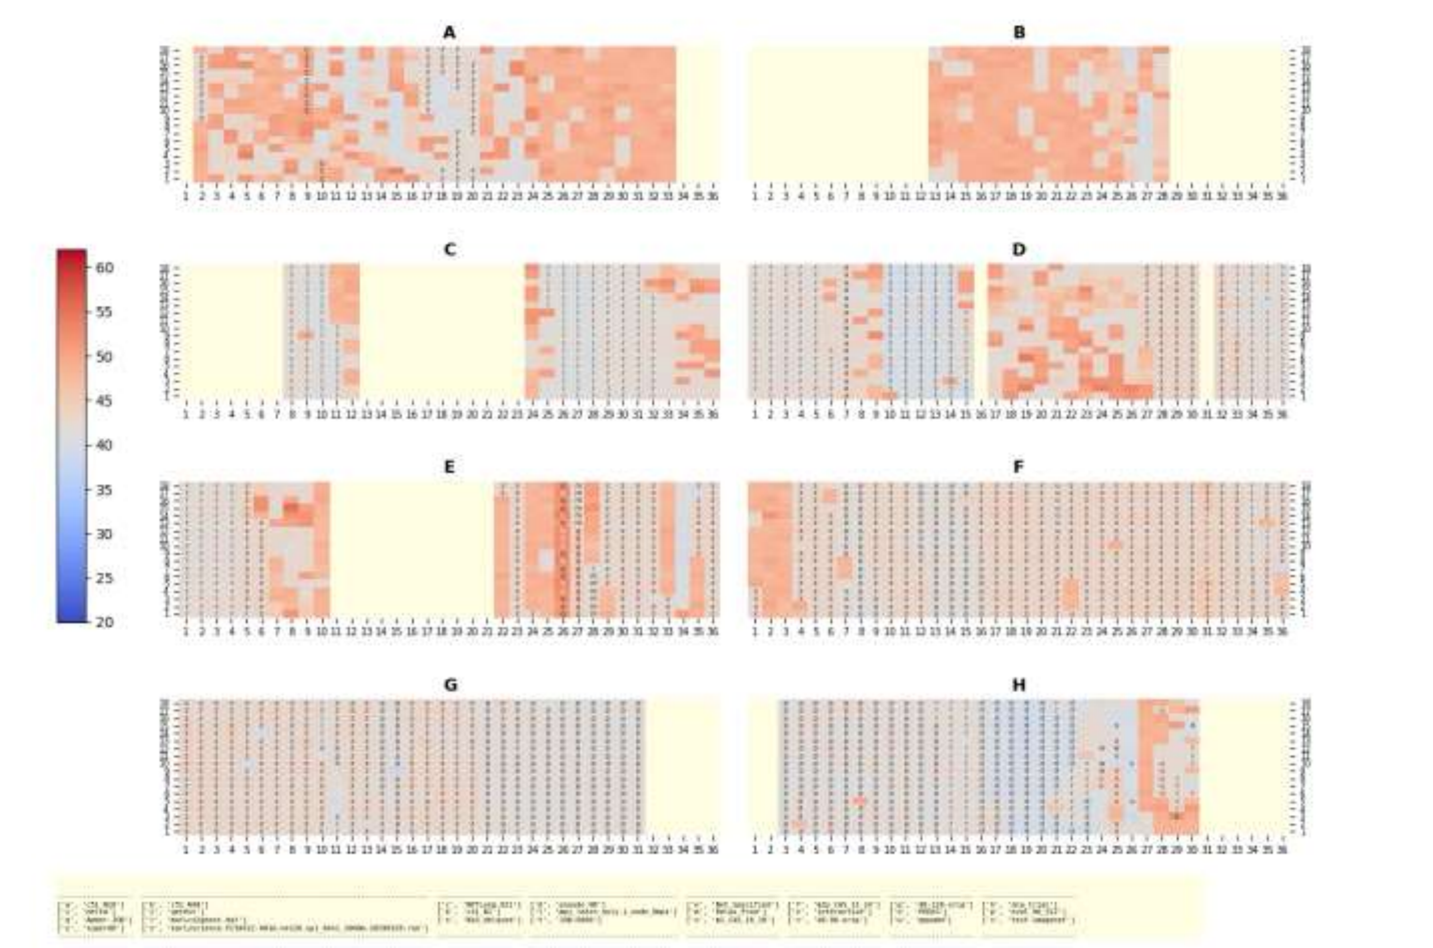

This Was Not Without Its Challenges…

Long ramp up time due to contract, legal & solidifying SOW requirements.

Resource constraints on both sides due to scheduling and unforeseen delays.

Had to work within the confines of the Summit security model – no direct access - containers saved the day.

Limited computer resources due to h/w budget. Built everything shown on 3 nodes.

But The Effort Was Worth It.

Overlay additional disparate data sources in a production environment that is both scalable and elastic.

Provide predictive analytics for water temperature control based on scheduled jobs in the queue that can provide input to the PLCs.

Develop predictive algorithms that intelligently affect power/temp across multiple dimensions: CPU, GPU, memory, power.

Determine how this platoform can be utilized in Exascale systems.

Provide interoperability between expert systems and components to create self-healing, scaling and operational control.

Research how specific codes/algorithms affect job performance.


summit

Use Case #2

Aerospace Defense Contractor

Problem Statement

Client wanted understanding on cloud technologies and how to move certain workloads into the cloud

They had minimal cloud experience

Wanted on-site, hands-on training

Comparative analysis between major cloud providers

Alleviate concerns around security, data transfer, multi-tenancy, account management

Solution Provided

Designed a custom two day ‘Cloud 101’ course tailored to the project the client desired to migrate into the cloud

Engagement model included weekly consultation for ongoing feedback, technical direction and open dialog (“Are we doing the right thing?”)

Showcased best practices that can be used on all cloud platforms

Presented how to think about cloud services as simplified tasks vs single monolithic HPC jobs

CLIENT PROPOSED ARCHITECTURE

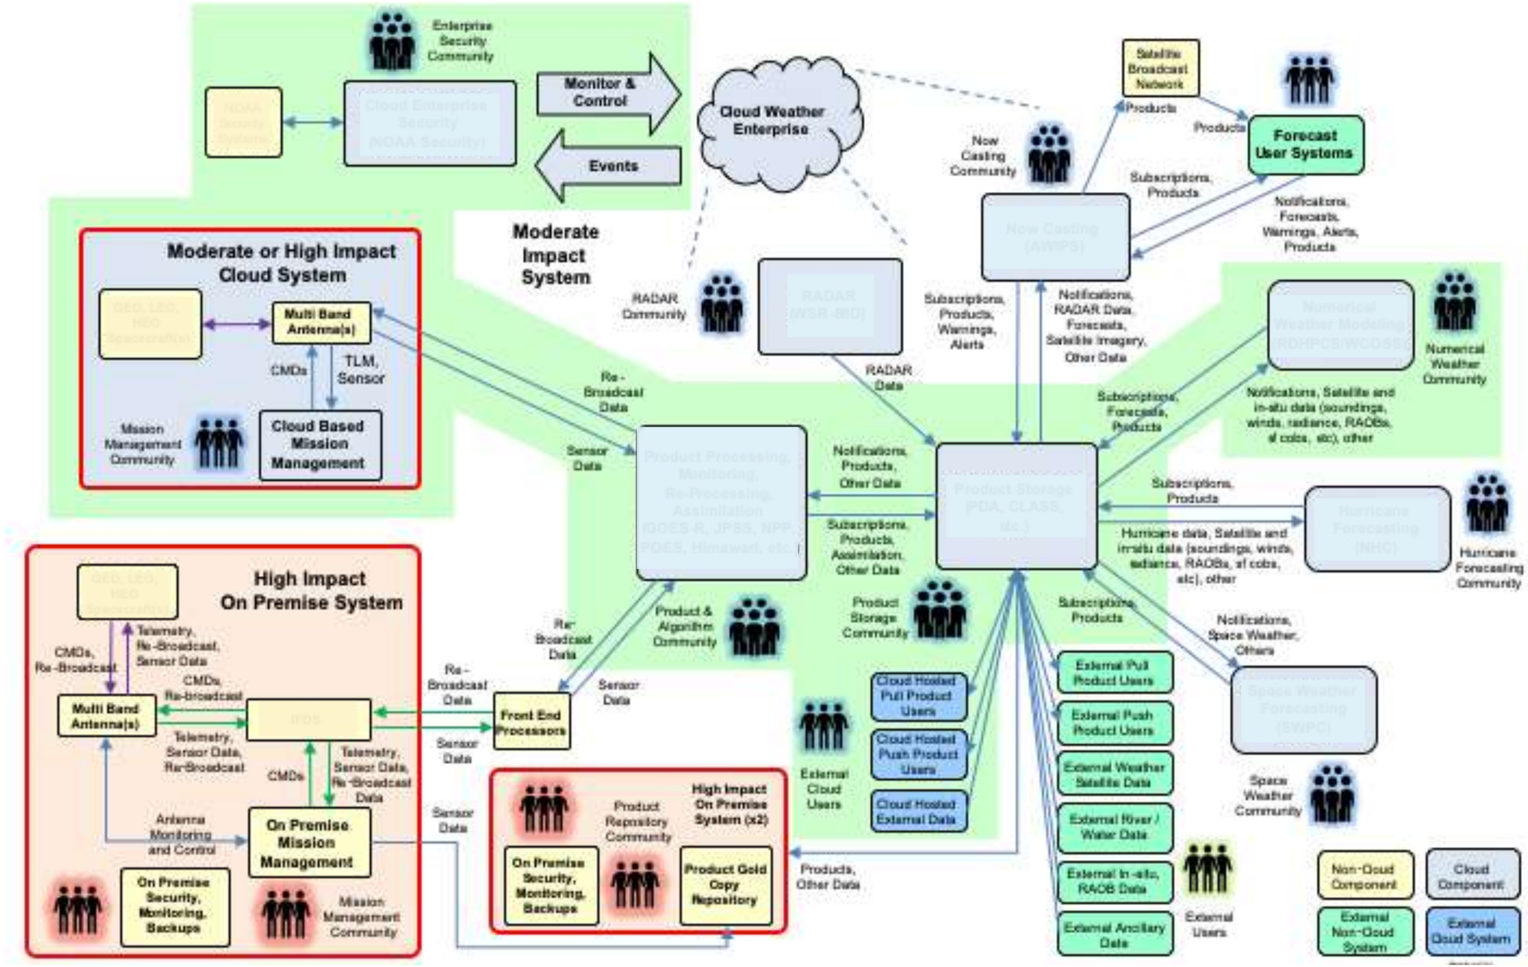

Client Solution Is…

Overly complex, both from a workflow and operational perspective.

Cannot meet availability requirements as structured.

Extremely difficult to scale, deploy and integrate new components.

On-prem / hybrid / cloud delineation is blurred.

Using mix of legacy, client/server Linux, HPC and standalone Windows.

PW CLOUD BASED ARCHITECTURE

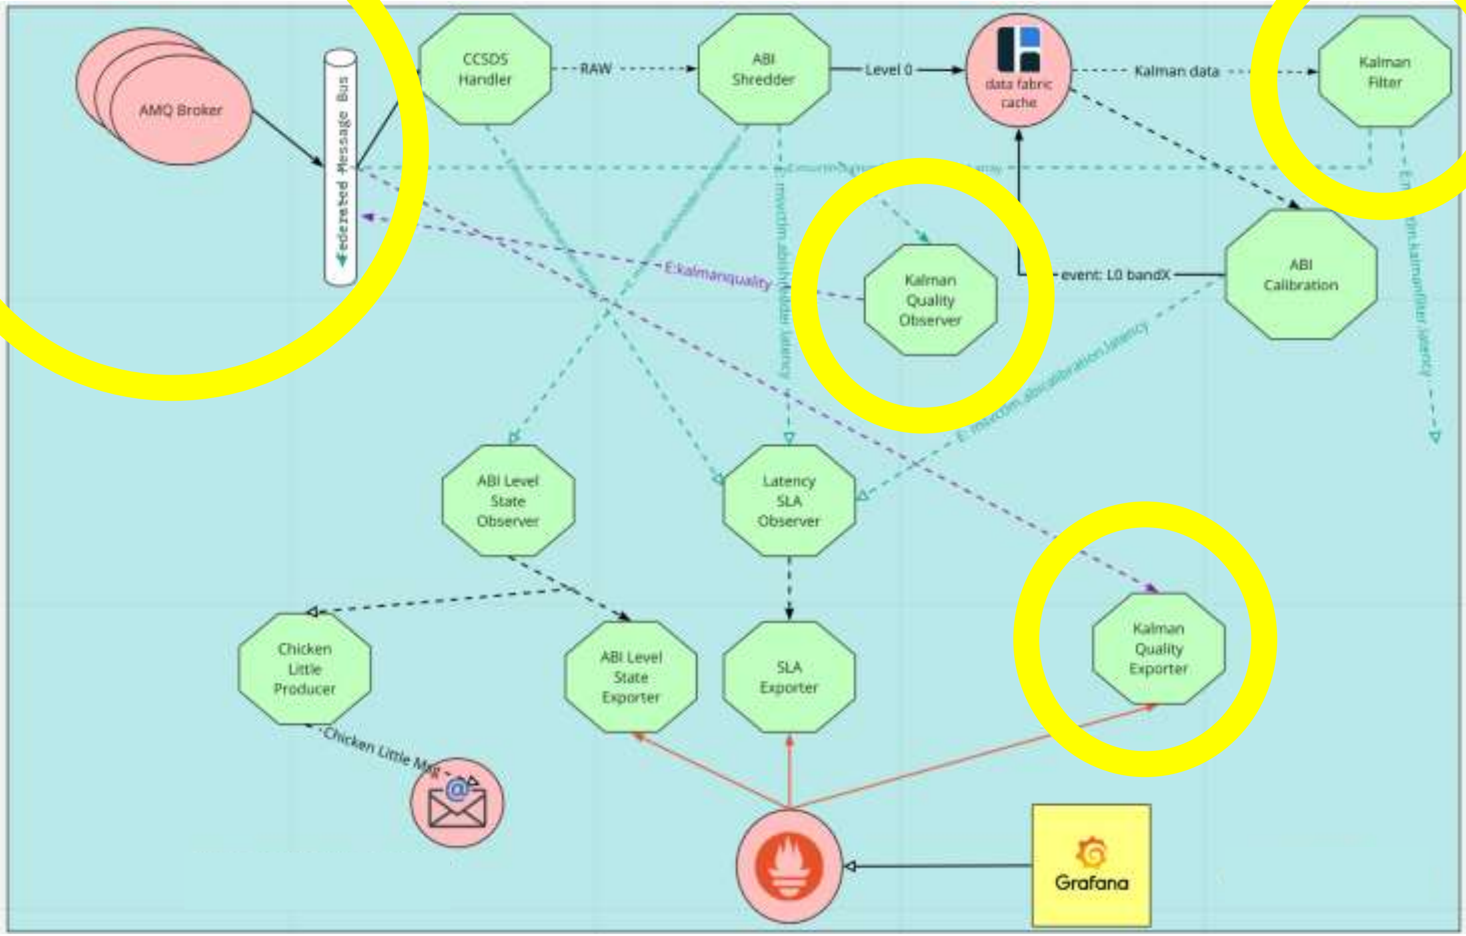

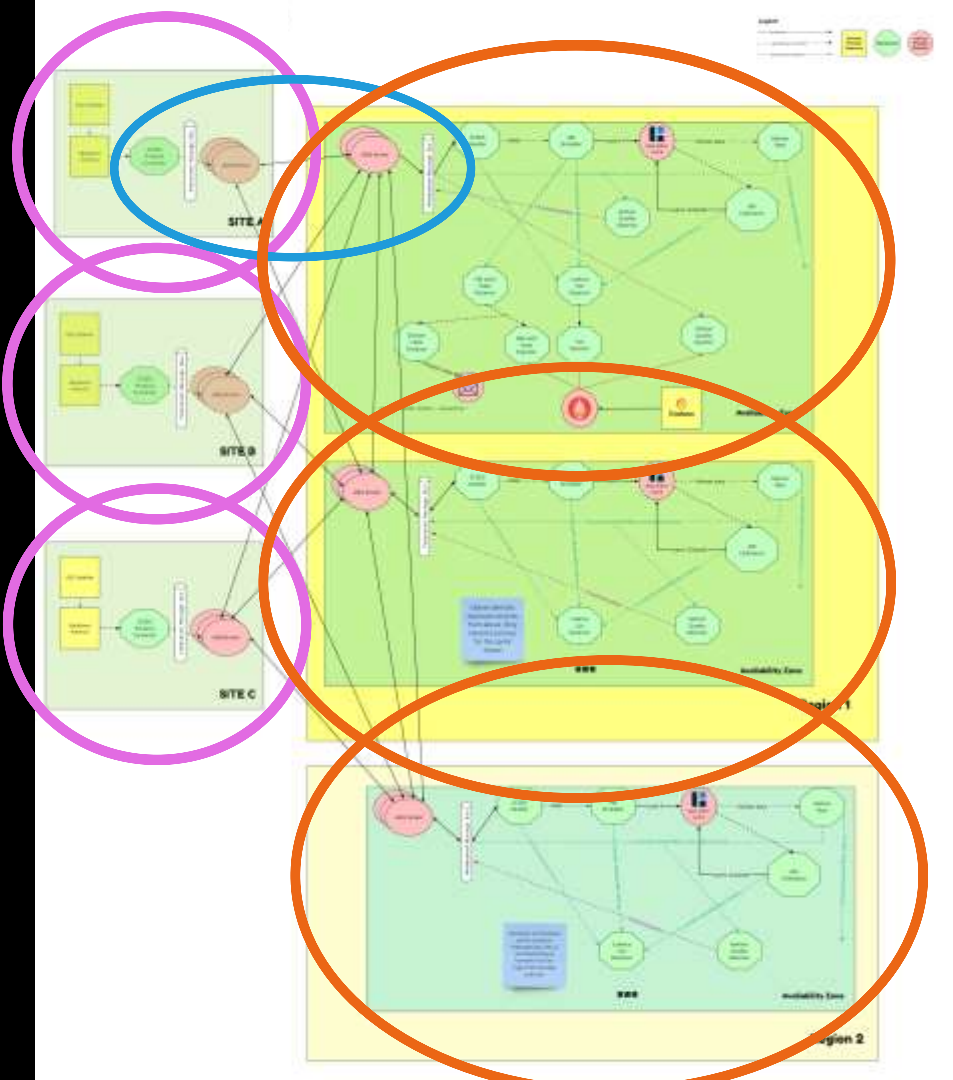
SITE B
SITE C

Cloud Based Solutions Applied

Applied a cloud-centric model leveraging modern architectures including containerized microservices

Now have availability zones across multiple regions for potential worldwide site deployments

Simple containerized deployment can be pushed to any cloud provider

Message bus allows faster and easier future integrations from disparate sites just by pushing container

On-prem / cloud integration delineated, but seamless

In Summary

Providentia Worldwide

HPC Integrations CAN Be Simplified

Leverage cloud based architectures that maximize efficiencies of scale.

Middleware is important - Use an ubiquitous transport system to deliver data, payloads, logs and control messages. Two use cases, two completely different reasons for bringing messages to bear.

Break large tasks into smaller ones to decrease complexity and introduce parallelism (where it makes sense).

Microservices can then be scaled independently as the needs arise.

Thank you.

arno.kolster@providentiaworldwide.com

@providentia_ww

www.providentiaworldwide.com