Neuromorphic Data Microscope CLSAC'16

October 28, 2016

David Follett
Founder, CEO
Lewis Rhodes Labs (LRL)
david@lewis-rhodes.com
978-273-0537

History

Neuroscience 1998 - 2012

Neuronal Spiking Models

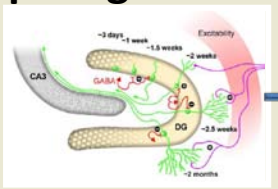

STPU

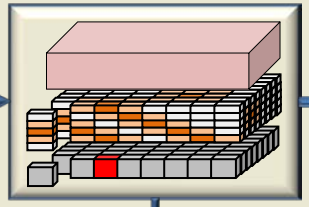

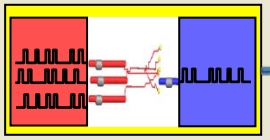

Neuronal Information Theory

NPU

Sandia LDRD 2013 - 2014

State Density $(n/m)^m$

Behavior Prediction

Pre-Attentive Processing

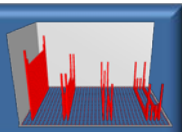

PCRE

$/Performance >100,000x
Behavioral Lexicon
Pervasive Analysis
Dynamic Reconfiguration

BECCA

$/Performance >100,000x
Power <100,000x
Weight <1,000,000x

HAANA 2015 - 2017

Neuromorphic Signal Processing
Image/video
Mixed Telemetry

Behavioral Analysis
Forensics
(aka *learning*)
Prediction

Neuromorphic Filters/Sensors
->Cyber<-
Big Data
Telemetry

Vision

Neuromorphic Processing Units (NPUs)

stunningly power efficient

at pattern matching

Data Center & Cloud Impact

profoundly changes economics

Why Neuromorphic?
Power/op & **Cost/op**

>500 IDS Servers

500,000 W
14,000 lbs
$2,500,000

All 4,036 "Emerging Threats" PCREs (1/16) Pervasive examination (no directives)

Performance Gains
Power/op: > 1M
Cost/op: > 1M

5 FPGA Neuromorphic Cyber Microscopes

2,500 W
200 lbs
$125,000

1 GbE **1 GbE**

40 GbE

1 ASIC Neuromorphic Cyber Microscope

LRL

50 W, 1 lb
$500

Neuromorphic is very Different

Legacy Von Neumann Architecture (CPU)

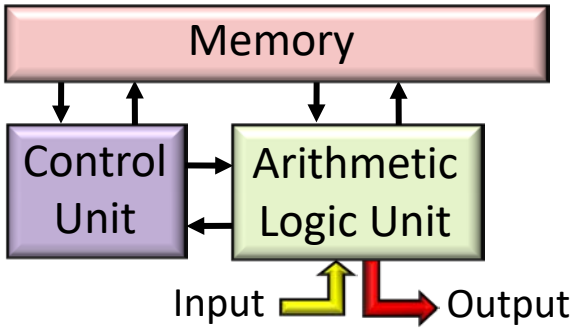

- **Complex processor**
  - Extraordinarily flexible
  - Data processing via sequential instructions
- **Simple memory**

Neuromorphic Processing Unit (NPU)

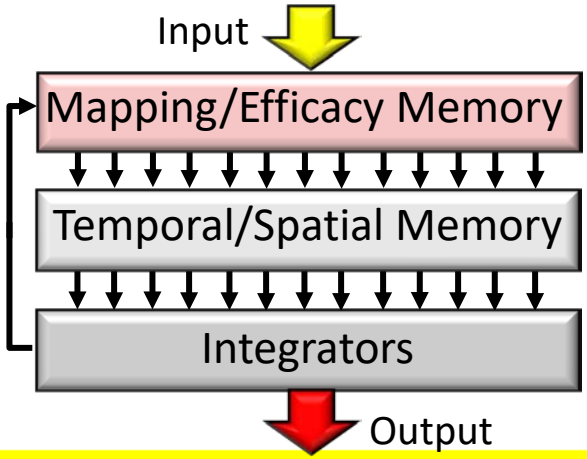

- **Simple processor**
  - Massively parallel integrators
- **Complex memory**
  - Data processing via efficacy & temporal/spatial mapping
  - Processing is multi-dimensional

Computer Science Perspective

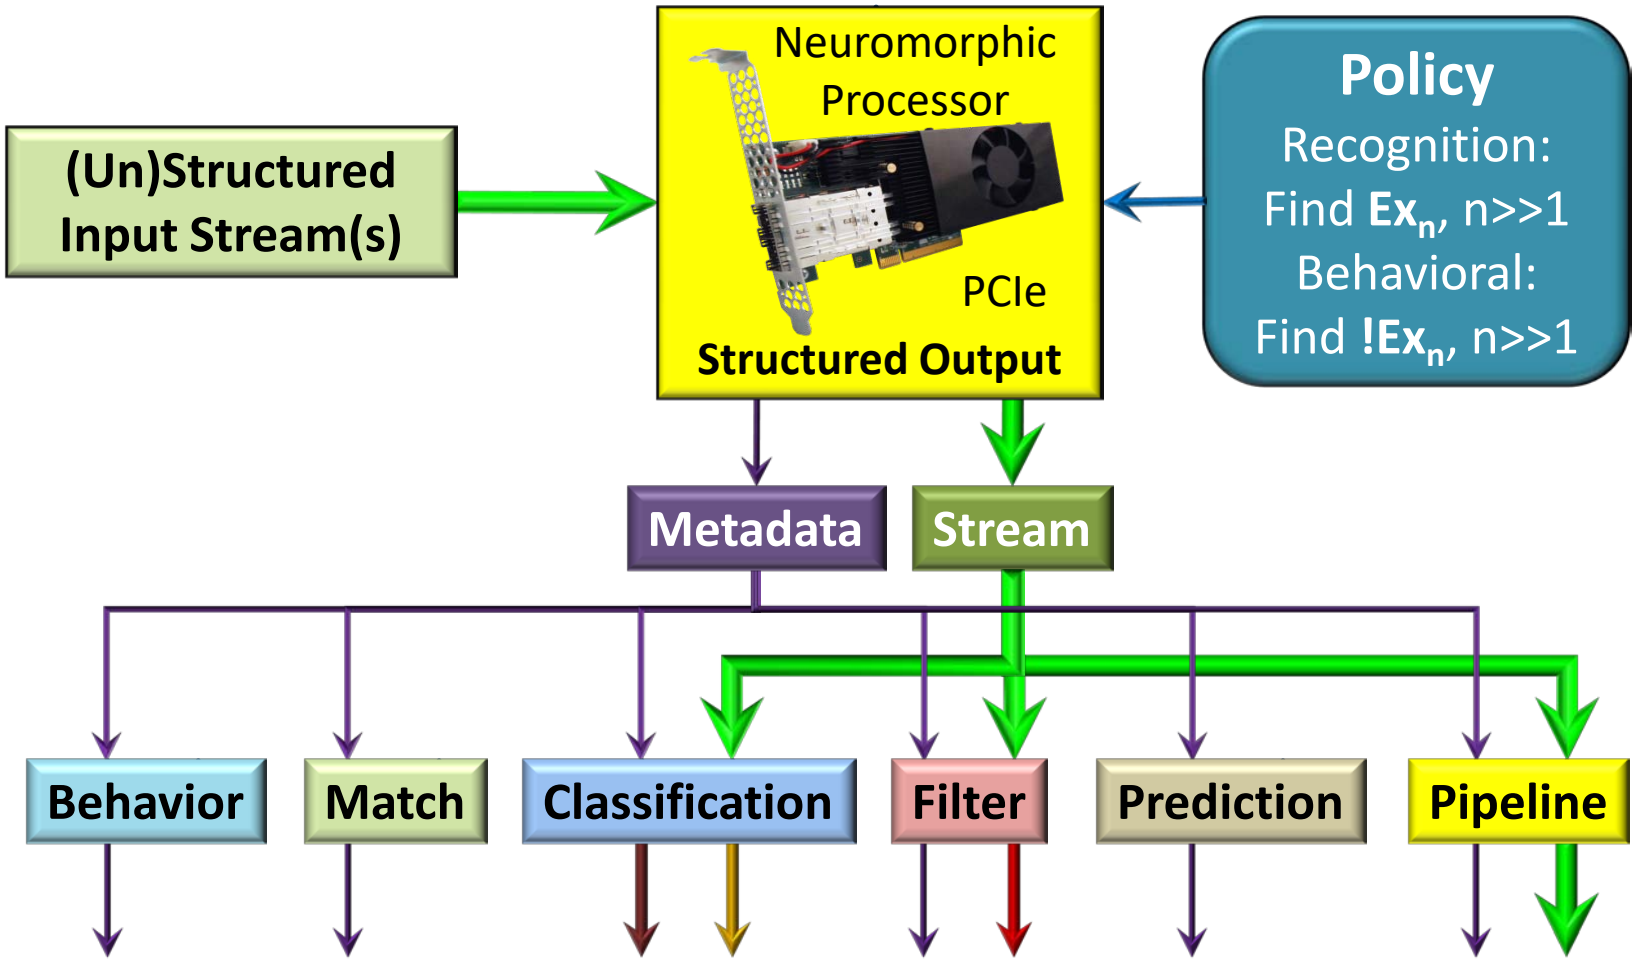

Some Interesting Features

- NPU integrates key mission requirements, ex.,
  - Context switching
  - Dynamic programmability
  - Behavioral characterization
  - Time & Order invariance
  - Pervasive analysis
  - Basic statistical operations
- Current device uses a single neuron type
  - Can extend HW architecture through novel neurons
  - Example: more complex statistical operations

Scalability

Device

Bandwidth x Expressions = Constant

FPGA

2.5 Gb/s x ≈ 1,000 Expressions
5 Gb/s x ≈ 500 Expressions
10 Gb/s x ≈250 Expressions
etc.

ASIC

20 Gb/s x ≈ 20,000 Expressions
40 Gb/s x ≈ 10,000 Expressions
80 Gb/s x ≈ 5,000 Expressions
etc.

System

Arbitrary Depth & Width

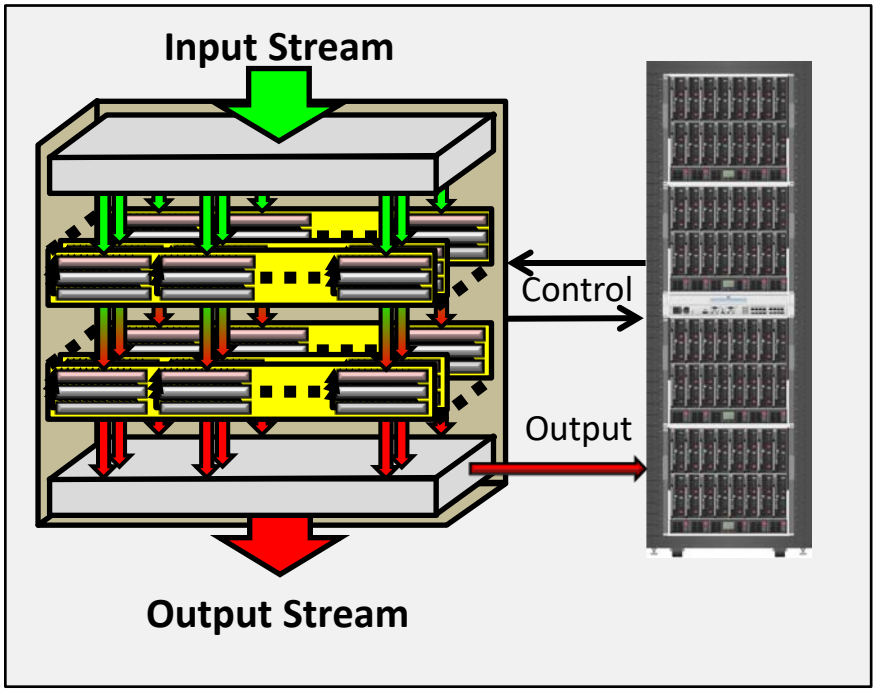

Latest Product

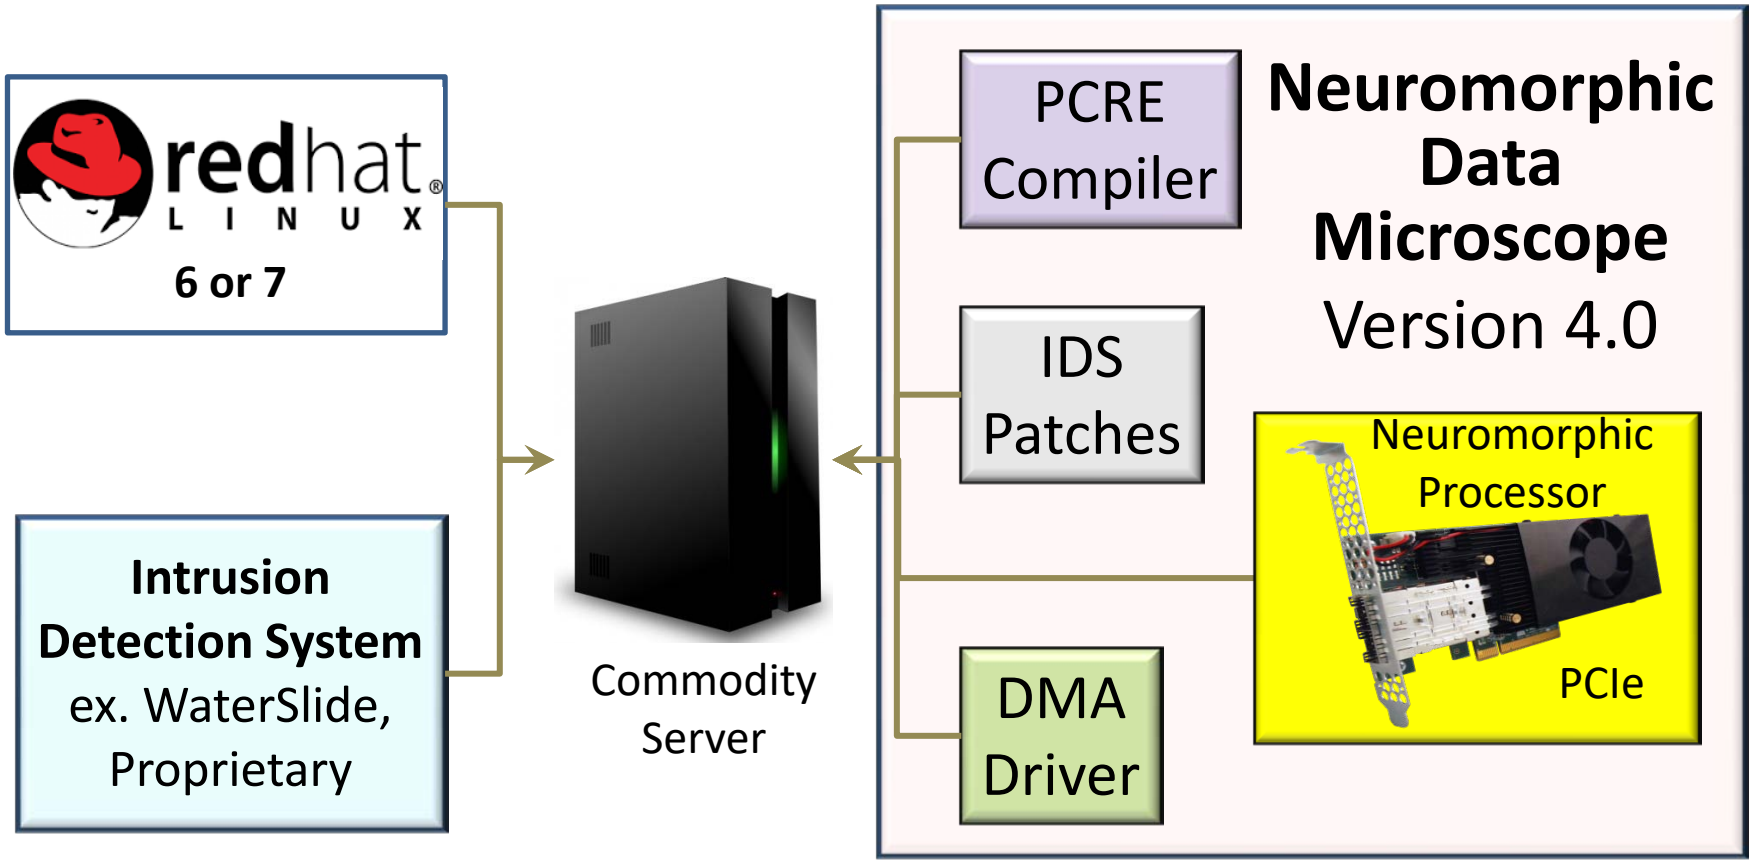

Standards Hide Complexity

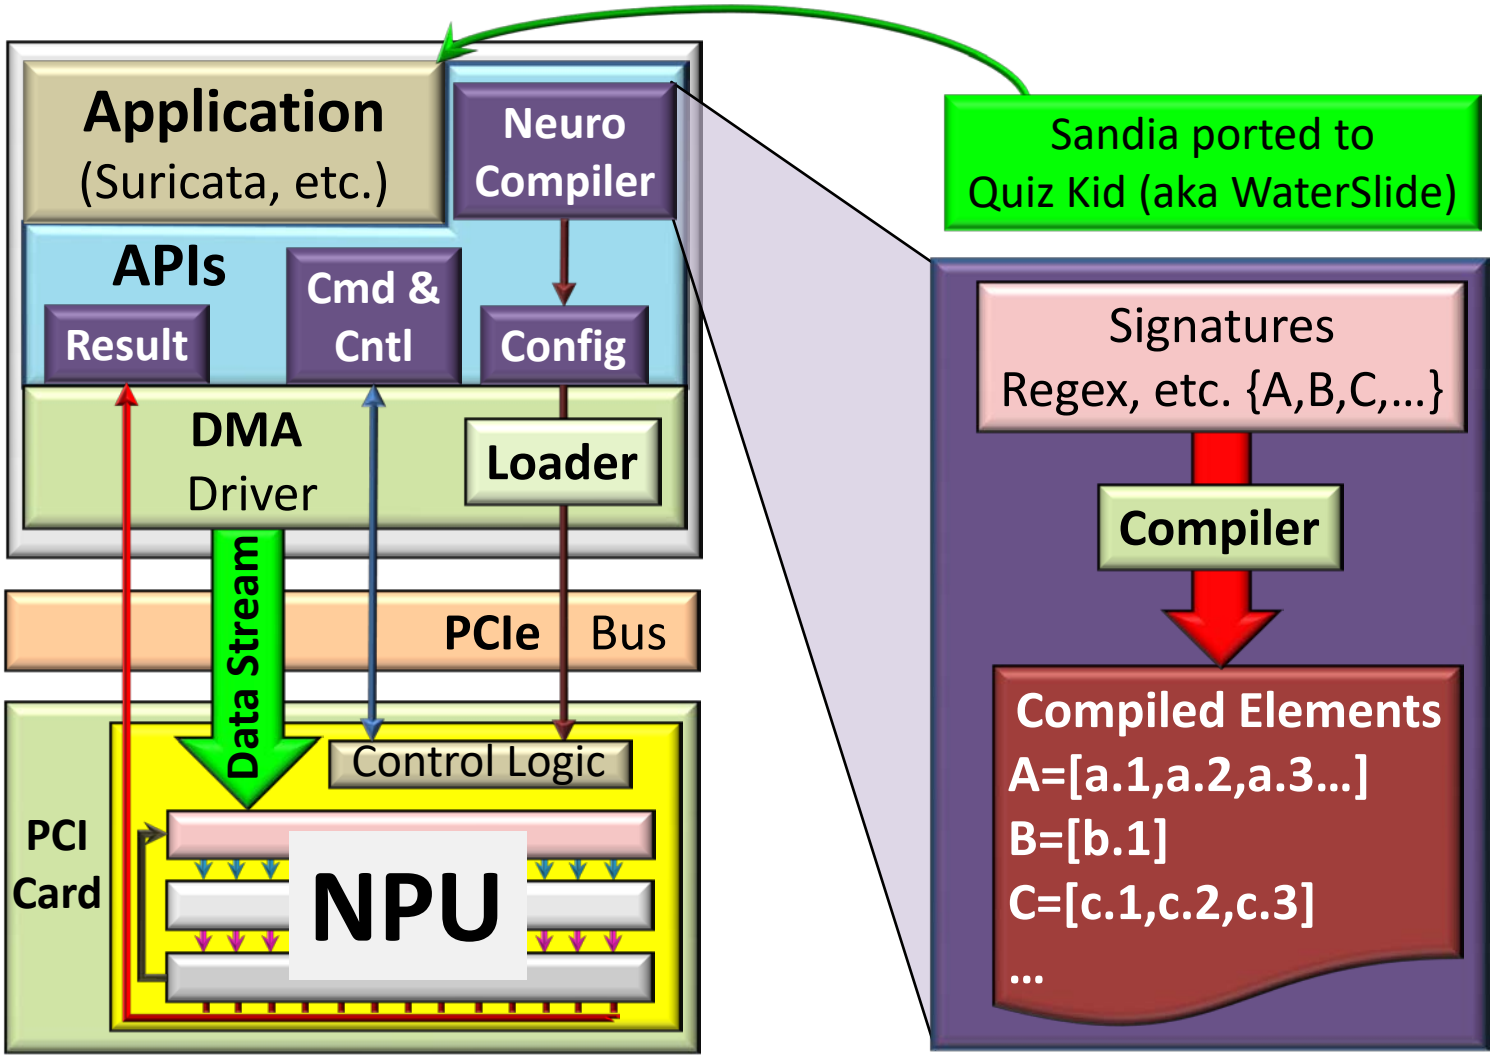

Throughput Efficiency

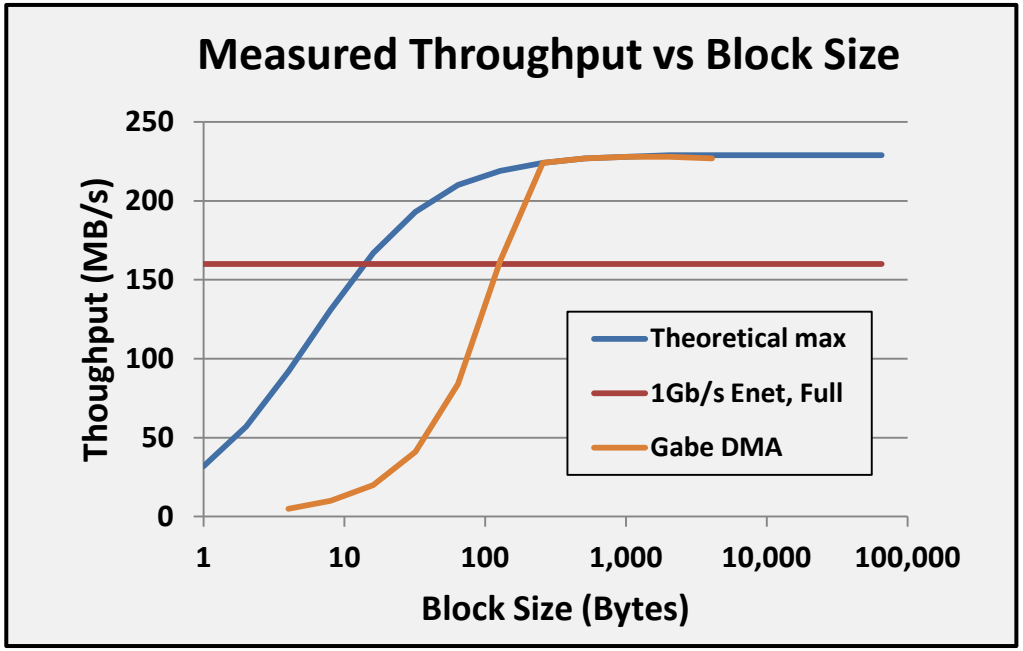

Note 1: Theoretical Max bounded by context switching
Note 2: Un-optimized generic Altera DMA Interface
Note 3: CUDA style DMA planed for next generation

Graph Analytics Scaling Example

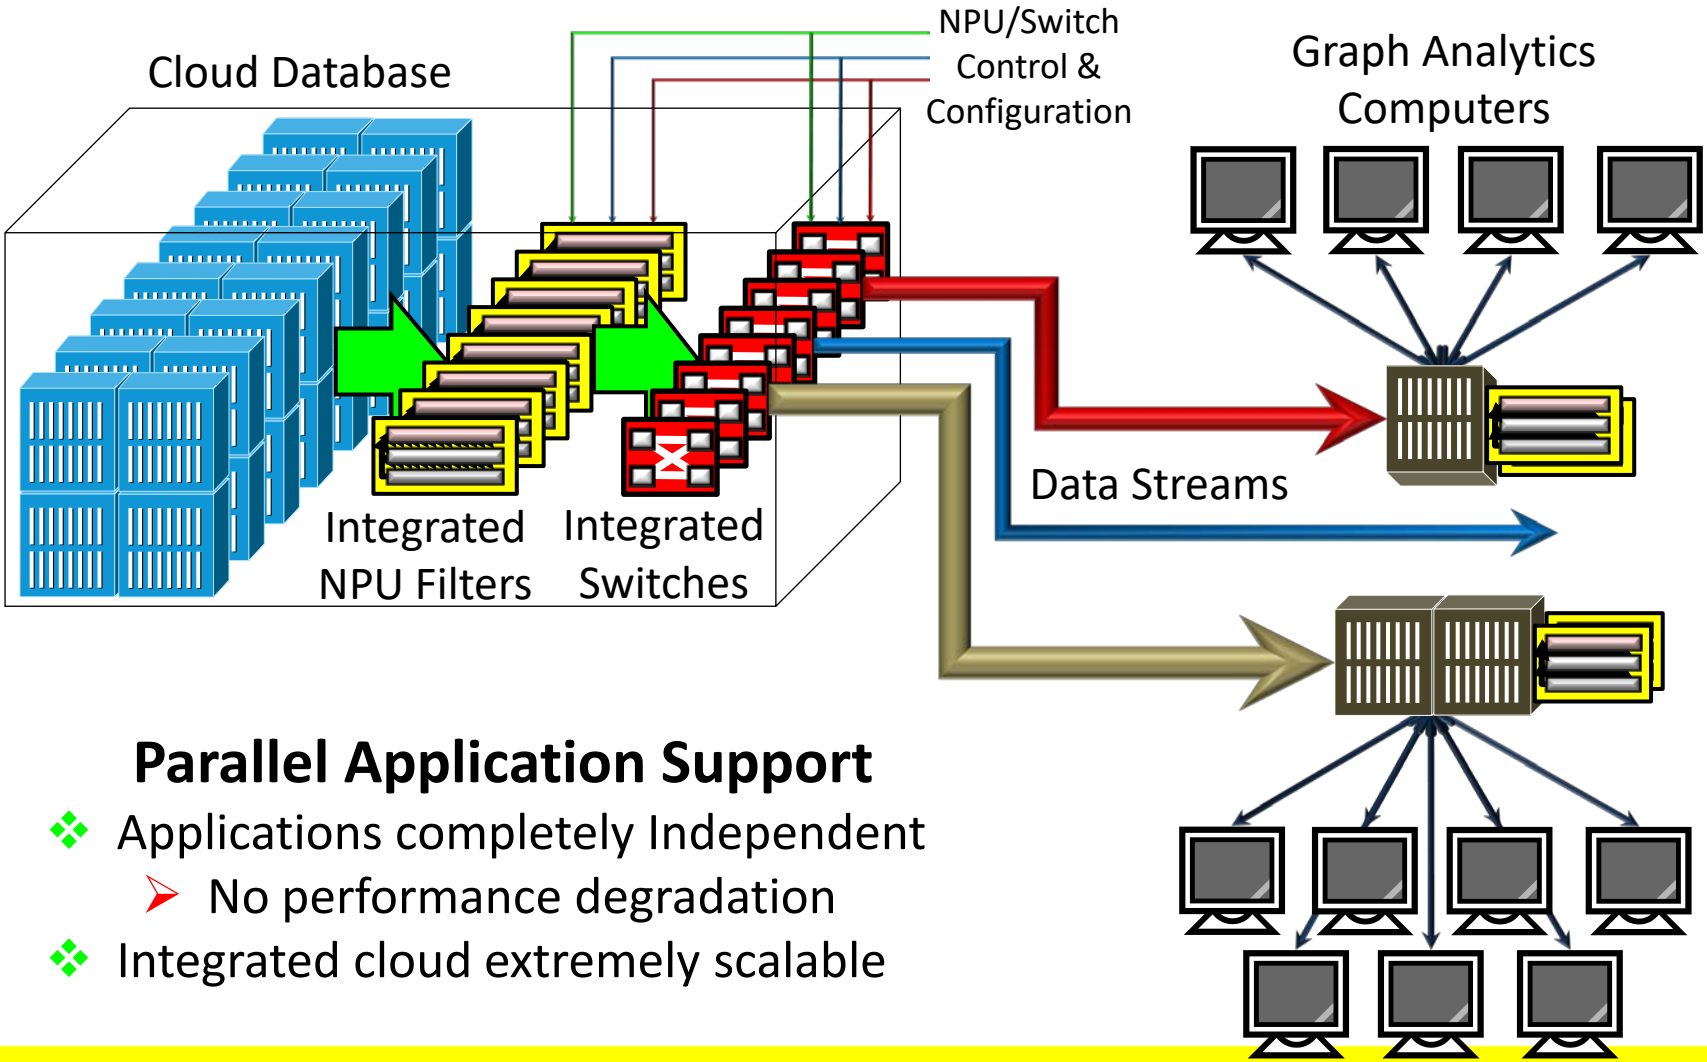

Parallel Application Support

- Applications completely Independent
  - No performance degradation
- Integrated cloud extremely scalable

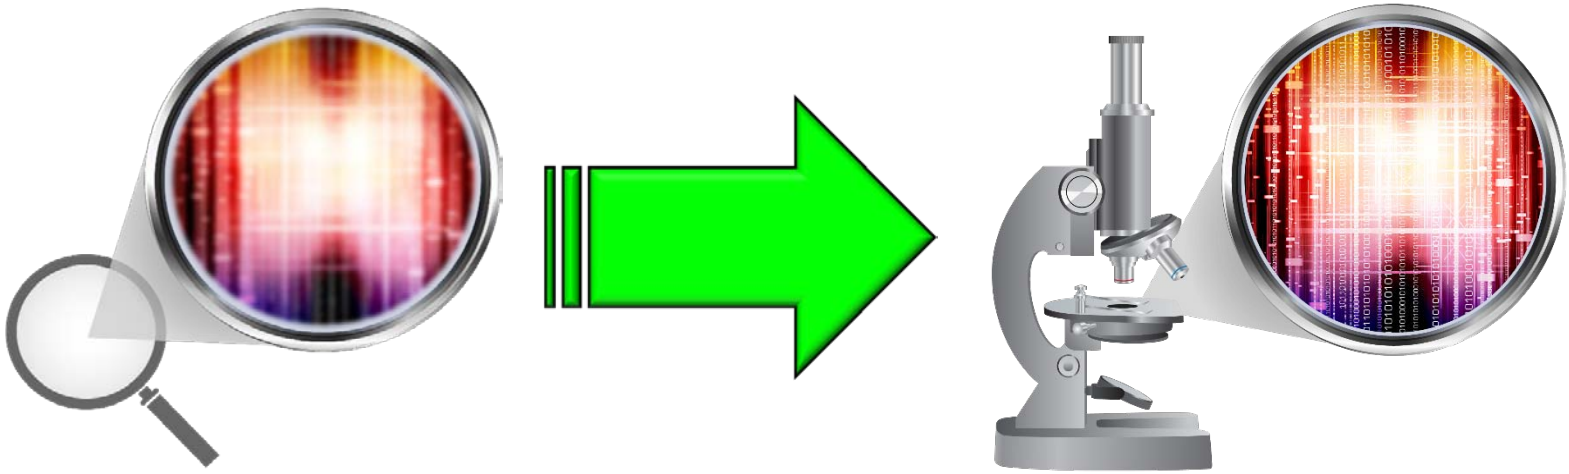

CYBER APPLICATION

Exemplar Intel Community IDS

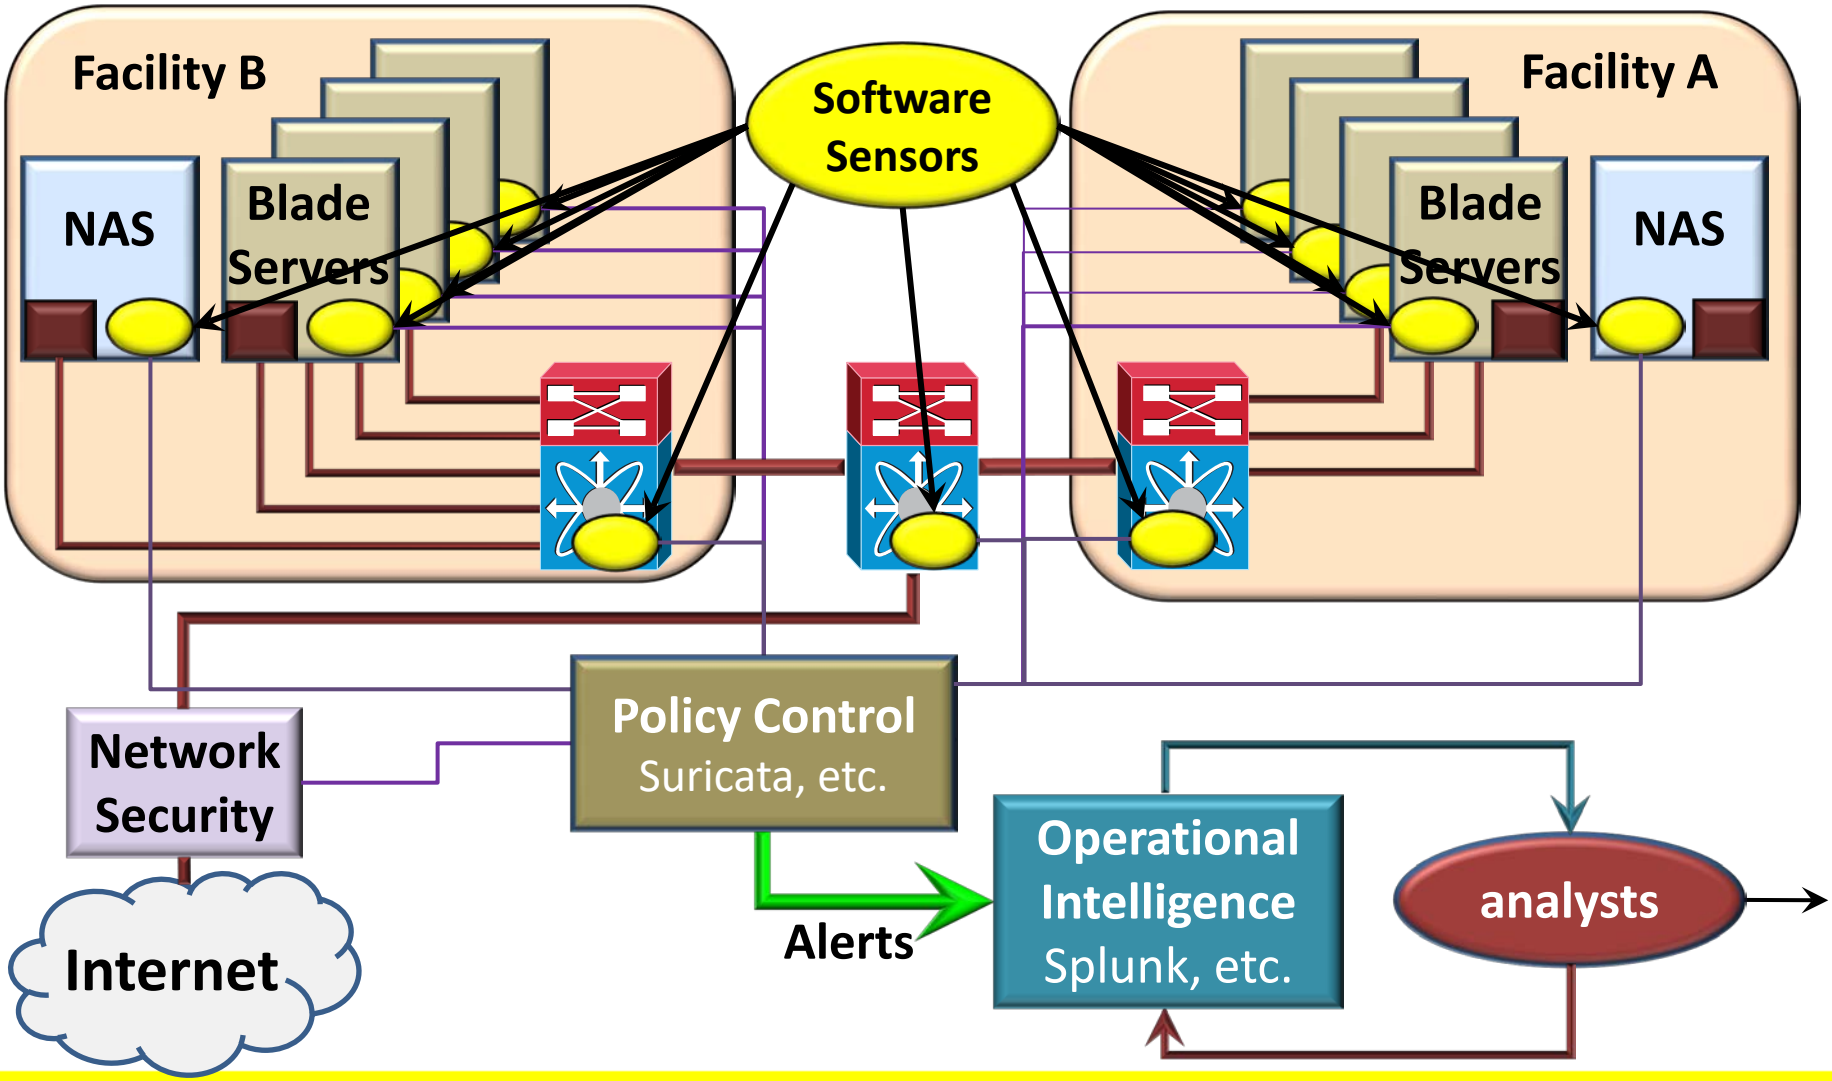

Practical Considerations

- Hardware sensor cost extremely high
  - ex. 10GbE IDS >$100k
  - Cost limits number & resolution of HW sensors
- Software sensors often resource intensive
  - ex. ROP detectors require most of the CPU
  - Cost limits number & resolution of SW sensors
- Analyst's priority, reduce False Negatives
  - Achieved by detuning sensors, ie. large # of False Positives
  - Major source of noise, direct result of sensor cost
- Detuned sensors are more vulnerable to attack
  - Spoofing & Flooding are common

Analyst's Top Priority

Signal/Noise is killing analyst community

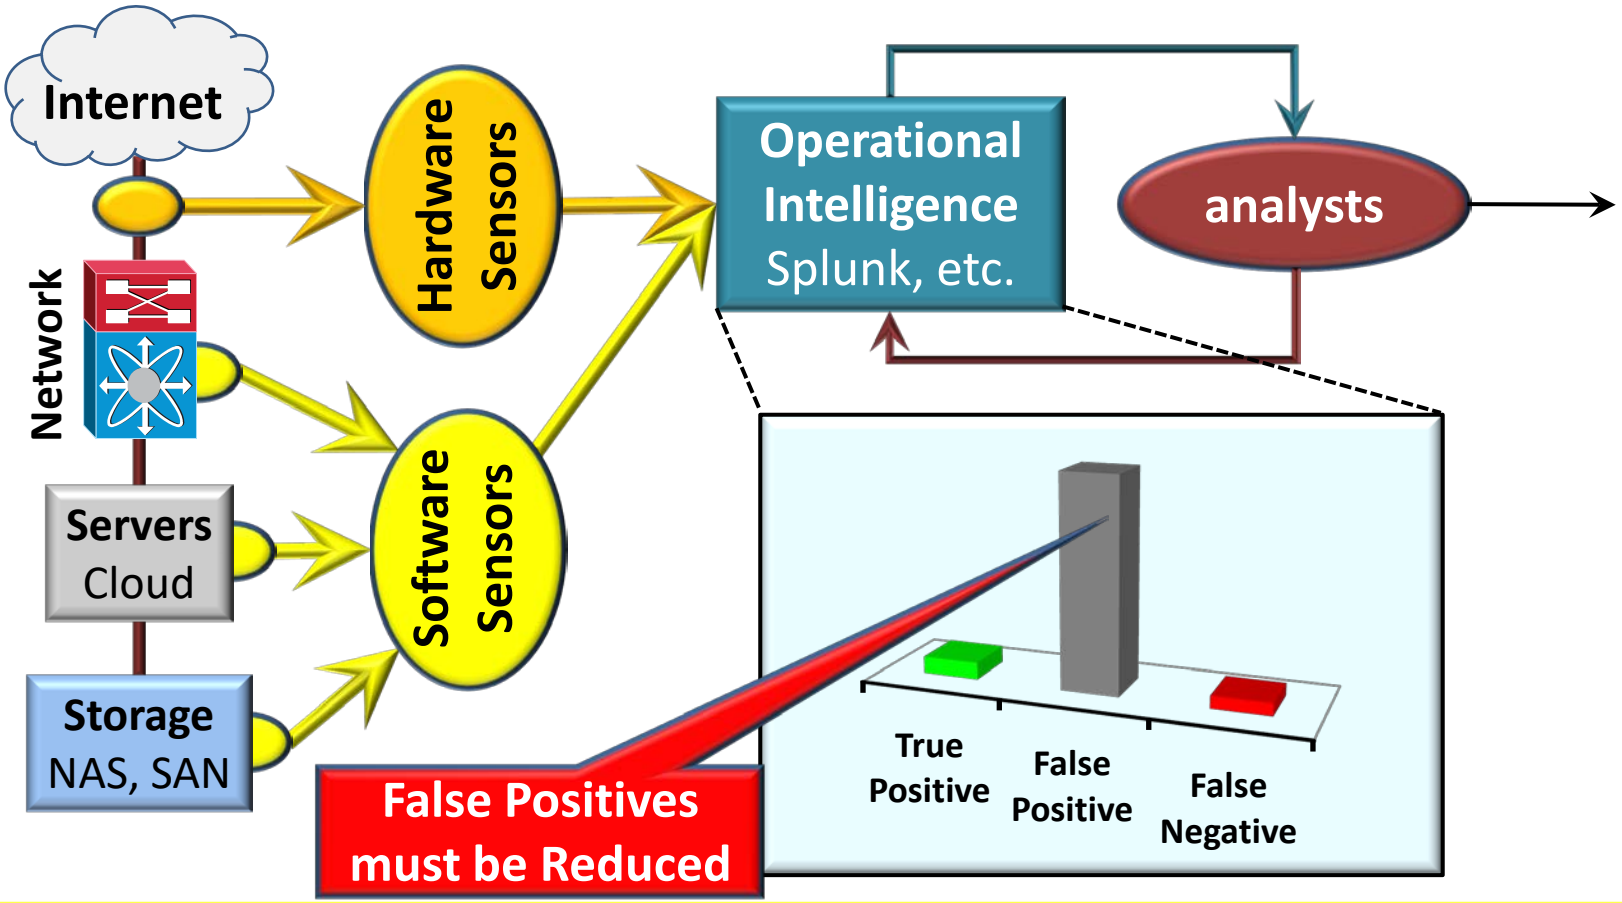

Root Cause: Resolution

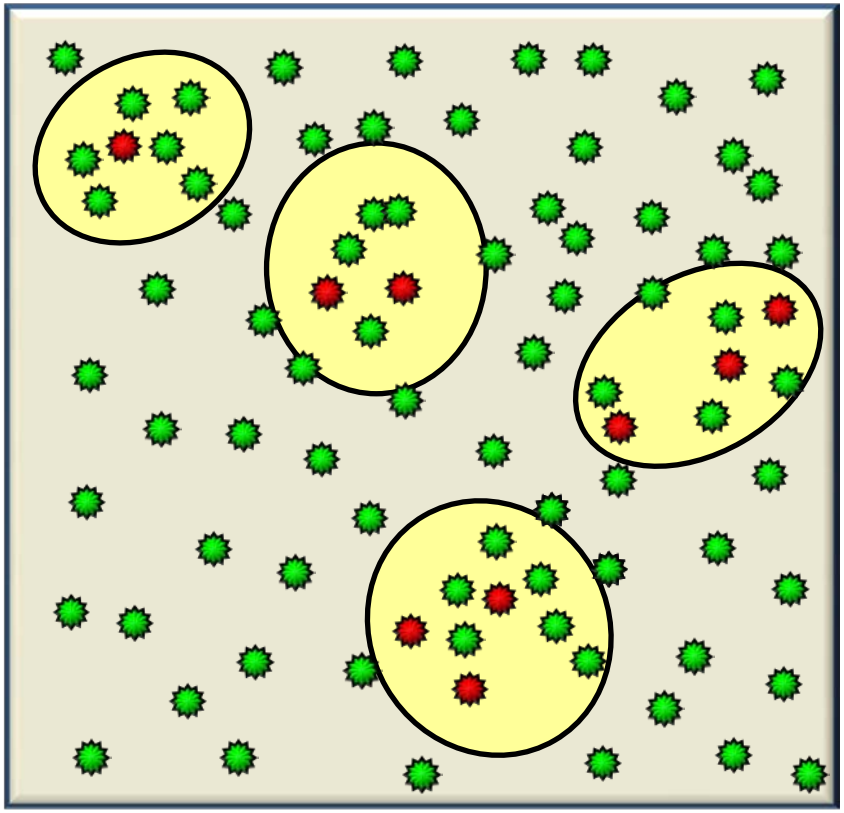

True Positives (**TP**)
Potential False Positives (**FP**)
Expression Coverage

State-of-the-Art Sensors

ex. **Suricata**

- **Cost** limits resolution
- **TP**s identified but
- Many **FP**s captured
- Splunk database,
  - Low Accuracy
  - Poor signal/noise ratio
  - It's still a haystack
- **S/N** is killing the analysts

Neuro: Resolution

Cyber Microscope

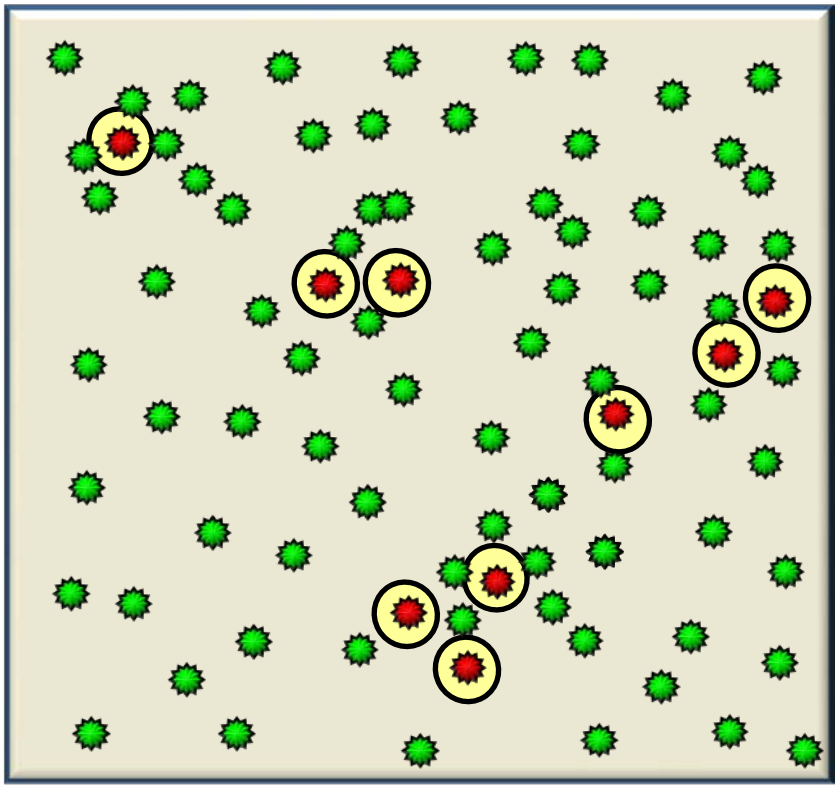

True Positives (**TP**)
Potential False Positives (**FP**)
Expression Coverage

Neuromorphic

- Speed creates resolution
  - Same number of **TP**s
  - Dramatically fewer **FP**s
- Greater Accuracy
- Higher Signal/Noise ratio
- Profound impact on analysts

Analyst's Second Priority

Reduce False Negatives

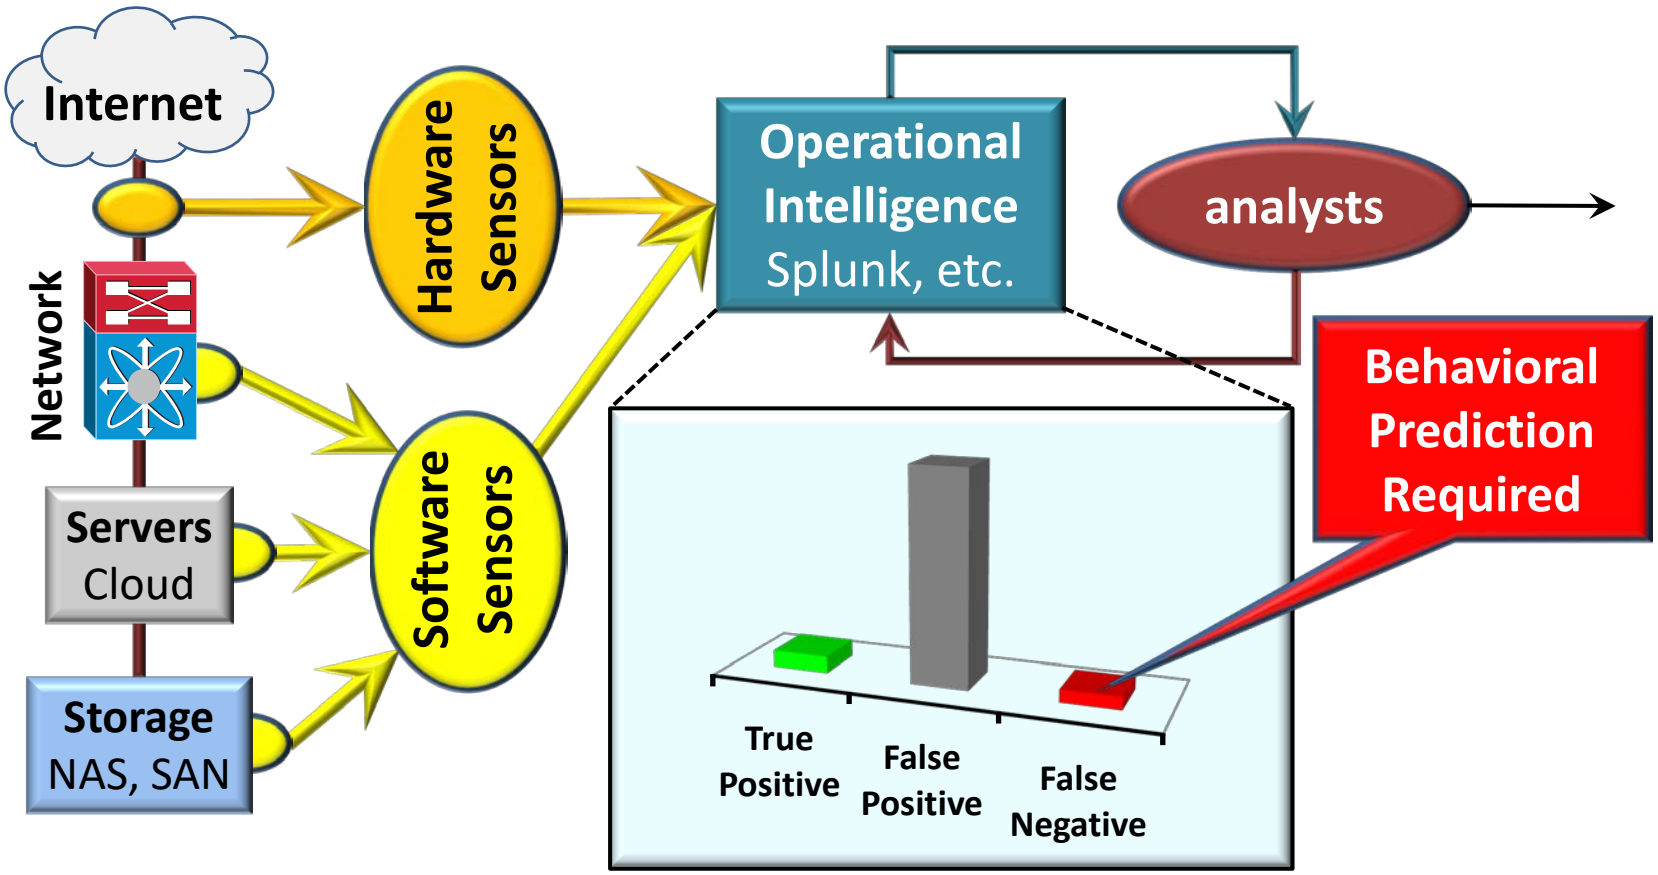

Root Cause: Temporal Variance

Simplest form of behavior prediction

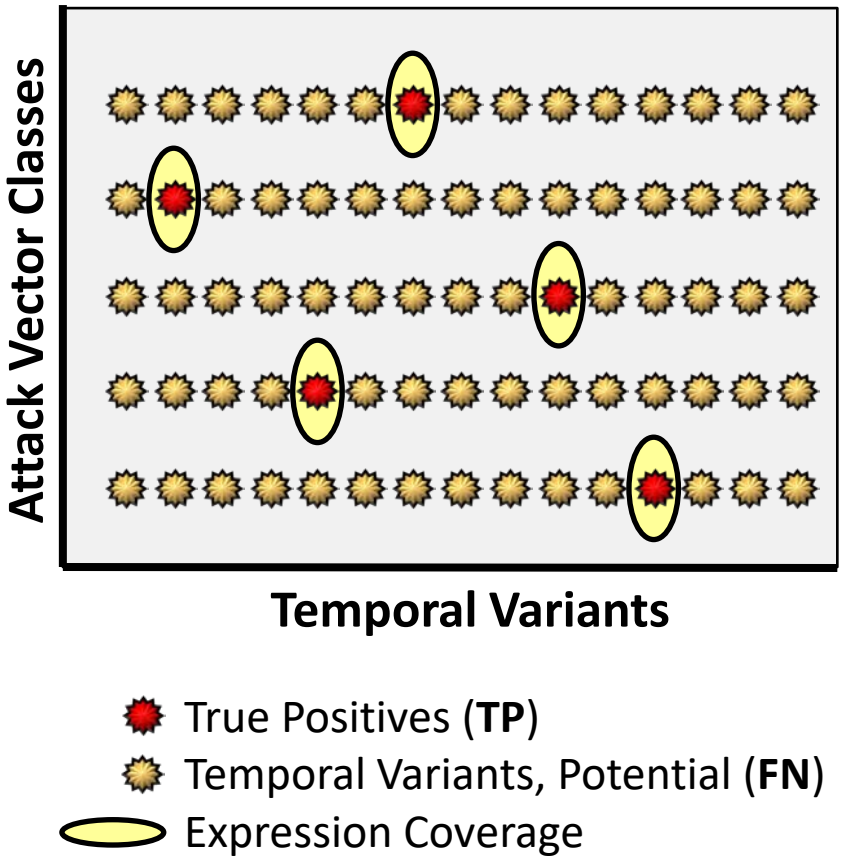

State-of-the-Art Sensors

ex. **Suricata**

- Temporal variance is common
  - Shifting offsets
  - Re-ordering
  - Easily implemented by attacker
- Very **costly** to address
  - Pervasive analysis
  - Associative analysis

Neuro: Temporal Variance

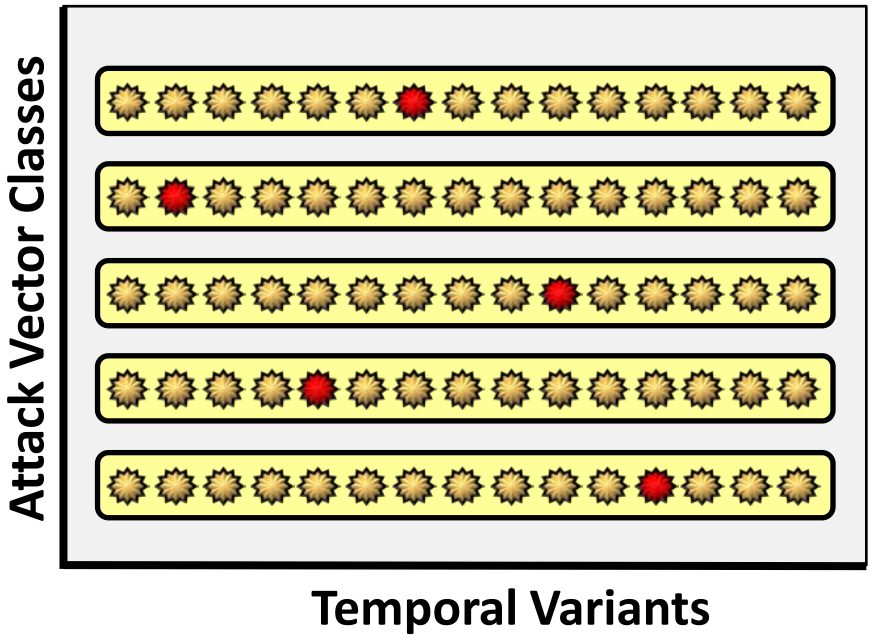

Neuromorphic

- Pervasive analysis is innate
  - Evaluates every byte
  - Limiting this costs resources
- Associative analysis is innate
  - Metadata reordering
- Reduced False Negatives **FN**
  - Behavioral Prediction
- Profound impact on Analysts

Neuro Addresses Core Issues

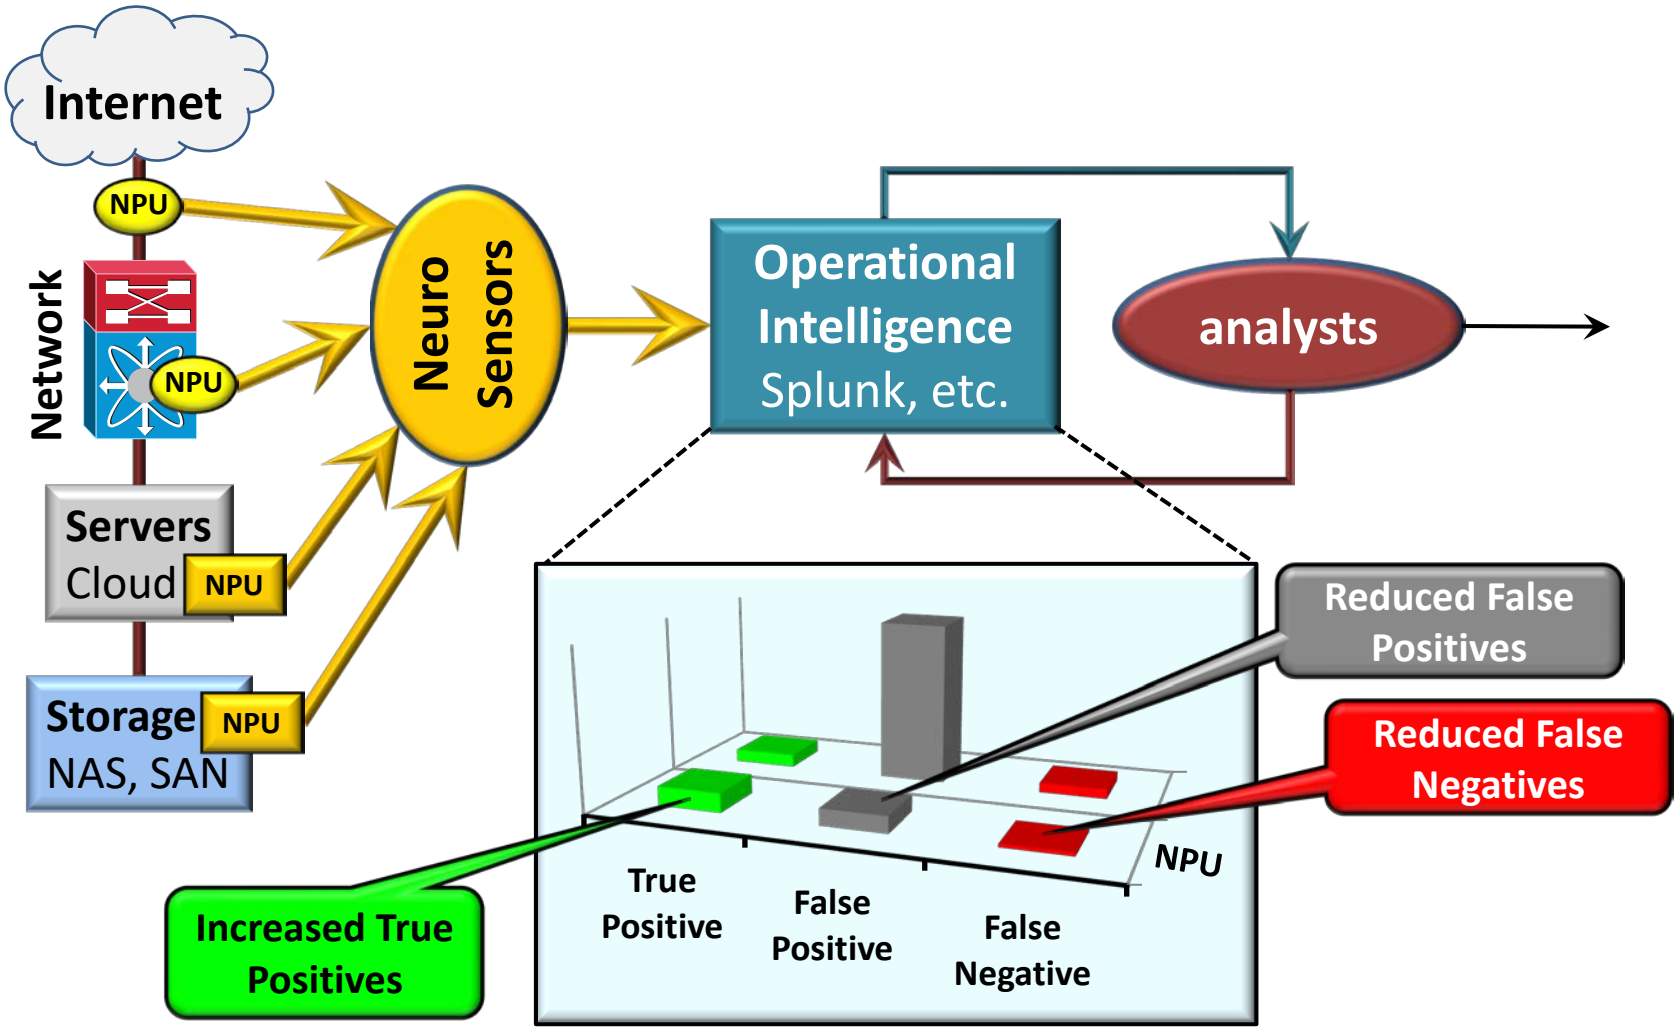

Operational Control

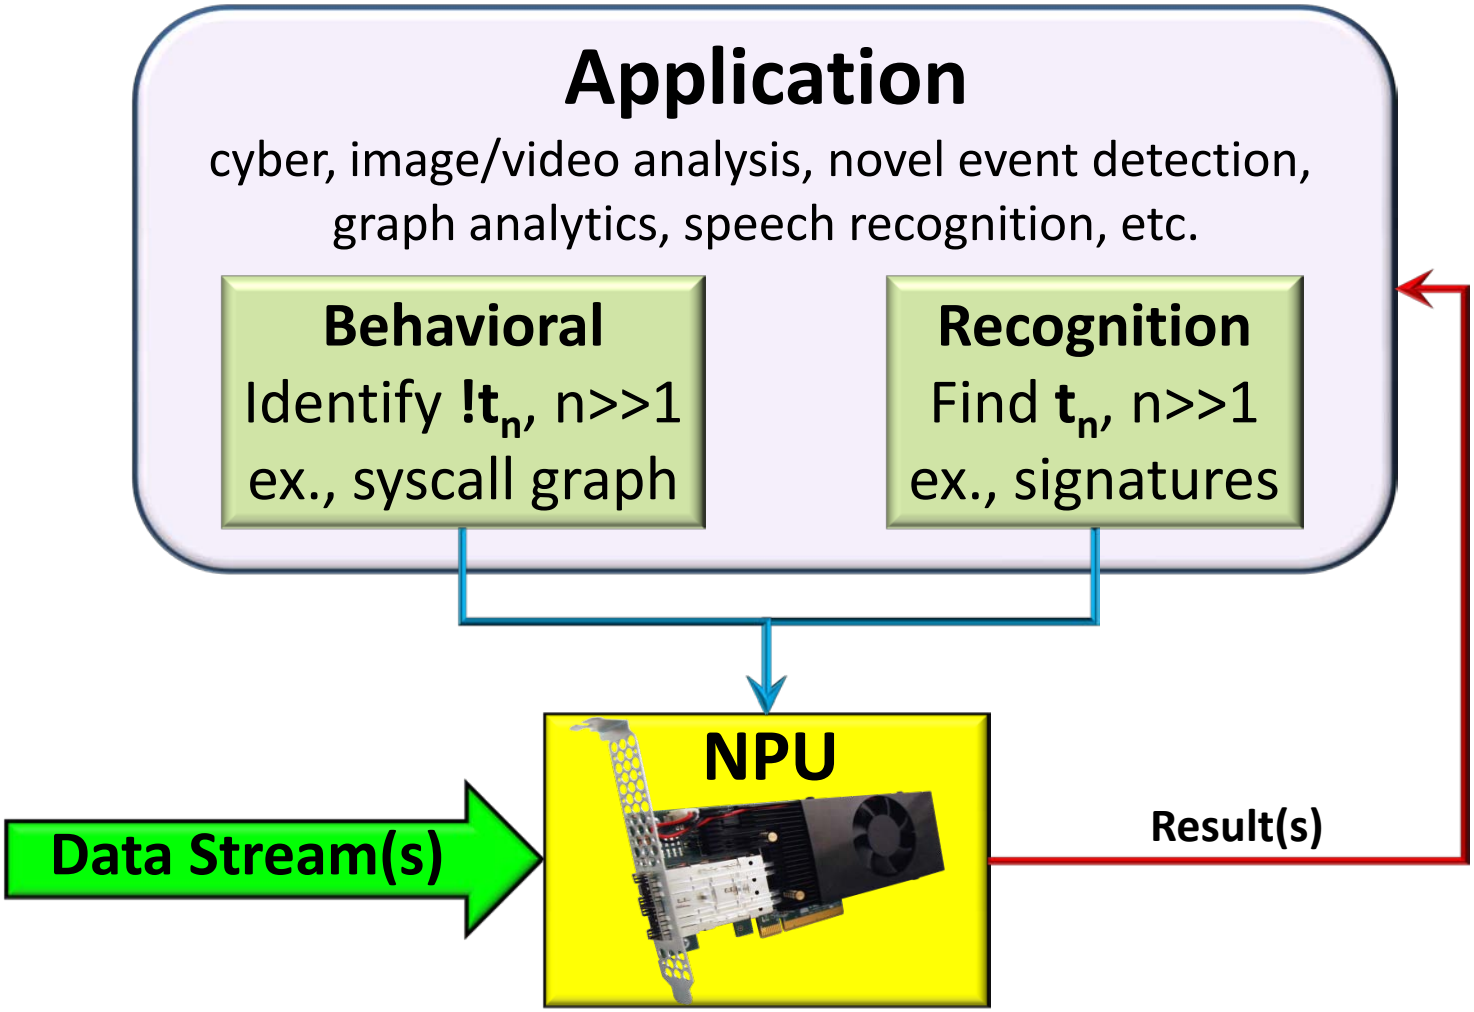

IC Analyst's Vision

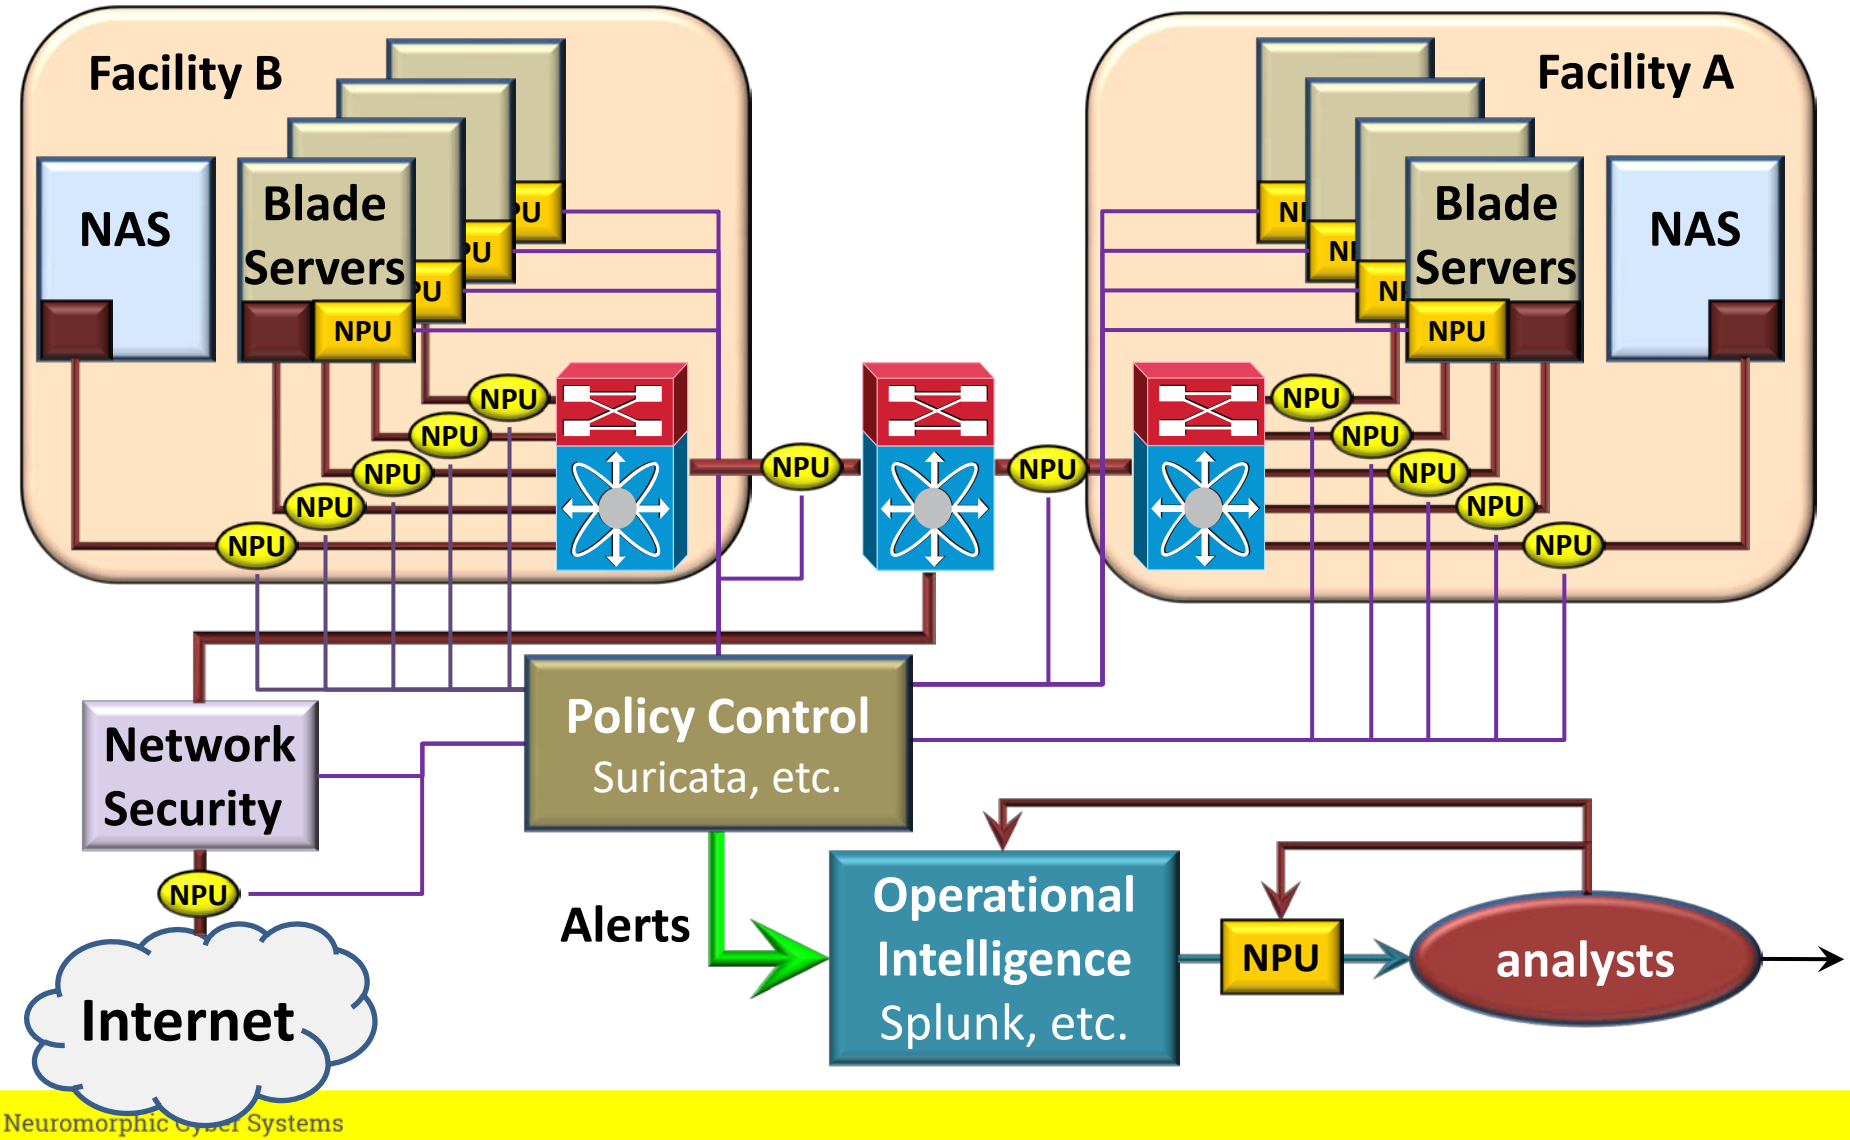

Cyber Microscope Product Rollout

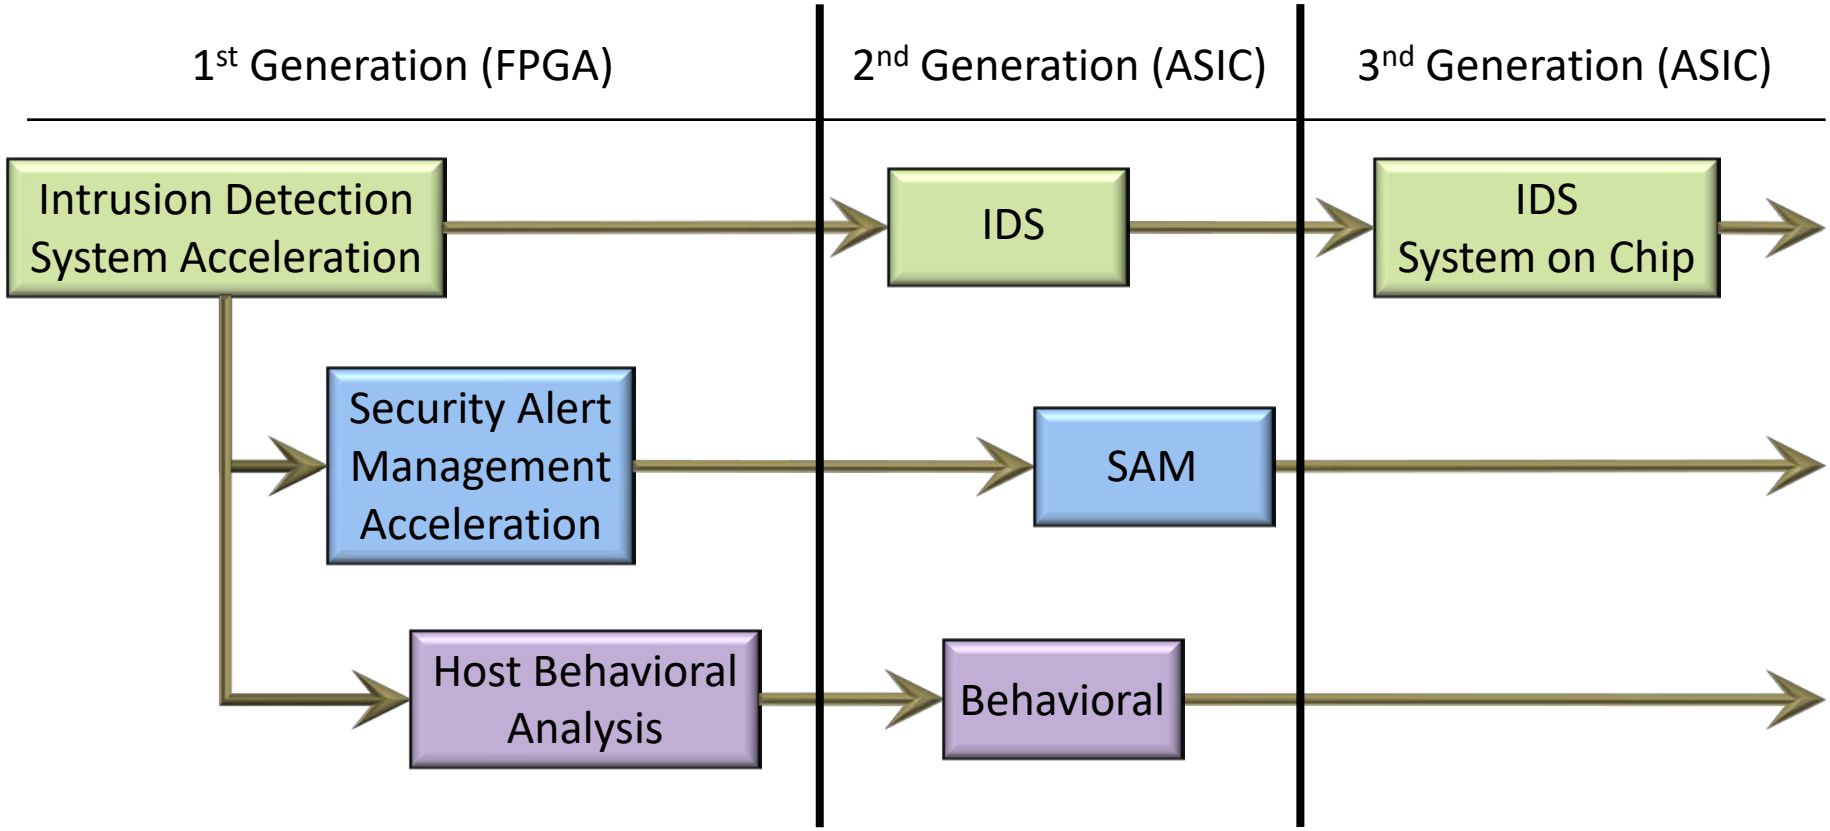

Conclusions

- Neuromorphic will revolutionize cyber defense
  - Dramatic reductions in power/op
    - FPGA, **>1,000x**
    - ASIC, **>1,000,000x**
  - Plethora of powerful novel features
    - Order & time invariant, Sessionization, Behavioral prediction
- Operational readiness is close
  - Compatible with existing standards & infrastructure
    - Sandia ported Quiz Kid (aka WaterSlide), 4 week effort
  - 4rd gen FPGA systems available in November